

Land Warfare and Complexity, Part II: An Assessment of the Applicability of Nonlinear Dynamic and Complex Systems Theory to the Study of Land Warfare (U)

Andrew Ilachinski

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited


Center for Naval Analyses

4401 Ford Avenue • Alexandria, Virginia 22302-1498

DTIC QUALITY INSPECTED 2

19990429 014

Approved for distribution:

A handwritten signature in black ink, appearing to read "Dr. Igor Mikolic-Torreira", written over a horizontal line.

Dr. Igor Mikolic-Torreira
Director, Systems and Tactics Team
Operating Forces Division

This document represents the best opinion of CNA at the time of issue.
It does not necessarily represent the opinion of the Department of the Navy.

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

For copies of this document, call the CNA Document Control and Distribution Section (703)


REPORT DOCUMENTATION PAGEForm Approved
OMB No. 074-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE July 1996	3. REPORT TYPE AND DATES COVERED Final	
4. TITLE AND SUBTITLE Land Warfare and Complexity, Part II: An Assessment of the Applicability of Nonlinear Dynamics and Complex Systems Theory to the Study of Land Warfare			5. FUNDING NUMBERS C - N00014-96-D-0001	
6. AUTHOR(S) A Ilachinski				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Center for Naval Analyses 4401 Ford Avenue Alexandria, Virginia 22302-1498			8. PERFORMING ORGANIZATION REPORT NUMBER CRM 96-68	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Distribution unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 Words) The Commanding General, Marine Corps Combat Development Command (MCCDC) asked the Center for Naval Analyses to assess the general applicability of the new science to land warfare. "New Sciences" is a catch-all phrase that refers to the tools and methodologies used in nonlinear dynamics and complex systems theory to study physical dynamical systems exhibiting a "complicated dynamics." This report concludes that the concepts, ideas, theories, tools and general methodologies of nonlinear dynamics and complex systems theory show enormous, almost unlimited, potential for not just providing better solutions for certain existing problems of land combat, but for fundamentally altering our general understanding of the basic processes of war, at all levels. Indeed, the new sciences' greatest legacy may, in the end, prove to be not just a set of creative answers to old questions but and entirely new set of questions to be asked of what really happens on the battlefield. The central thesis of this paper is that land combat is a complex adaptive system. That is to say, that land combat is essentially a nonlinear dynamical system composed of many interacting semi-autonomous and hierarchically organized agents continuously adapting to a changing environment.				
14. SUBJECT TERMS computerized simulation, dynamics, land warfare, mathematical models, methodology, military science, modeling and simulation, technology, theory			15. NUMBER OF PAGES 162	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT	

"For the want of a nail the shoe was lost.
For the want of the shoe the horse was lost.
For the want of a horse the message was lost.
For the want of a message the battle was lost.
For the want of a battle the kingdom was lost.
And all for the want of a horseshoe nail."

How nonlinear dynamics and complex systems theory can make more precise the real meaning of this simple children's nursery rhyme, and how the ideas implicit in it can potentially be used to provide genuine insights into the fundamental processes of land warfare, is the subject of this report....



"Hey Dr. Operations Research, this complex systems theorist over here claims to have found a slightly simpler version..."¹

¹ A slightly altered form of a cartoon taken from J. Casti, *Complexification*, Harper Collins, 1994, page 10.

Table of Contents

Executive Summary	1
A Concise Summary of the Overall Assessment	1
Central Thesis	2
Eight Tiers of Applicability	3
Most Promising Applications	7
General Guidelines	9
General Discussion	11
Preliminary Remarks	11
The Charter	12
The Problem	13
The Approach	15
Two Preliminary Questions	16
<i>Question 1: What are Nonlinear Dynamics and Complex Systems Theory?</i>	16
<i>Question 2: Why <u>Land Warfare</u> and not Air Warfare or Naval Warfare?</i>	18
What Have We Learned From Complexity Theory?	19
What We Know	20
What We Don't Know	22
How is Work in the "New Sciences" Actually Done?	22
What are the Basic Questions to ask of a Complex System?	25
Patterns on a Battlefield	26
Combat as Soccer?	27
Models and Simulations: An Heuristic Discussion	29
What Makes a Model Useful?	30
Three Classes of Models	31
Computer Models	33
What Price Complexity?	34
Assessment of General Applicability	39
Risk vs. Potential Payoff	40
Tier I: General Metaphors for Complexity in War	43
What is a Metaphor?	43
Metaphors and War	45
Metaphor Shift	48
Tier II: Policy and General Guidelines for Strategy	54
What Does the New Metaphor Give Us?	54
Policy	54
Organizational Structure	55

Intelligence Analysis	57
Policy Exploitation of Characteristic Time Scales of Combat	58
Tier III: "Conventional" Warfare Models and Approaches	60
Testing for the Veracity of Conventional Models	60
Non-Monotonicities and Chaos	61
Minimalist Modeling	63
Generalizations of Lanchester's equations	64
Nonlinear dynamics and chaos in arms-race models	67
Tier IV: Description of the Complexity of Combat	68
Attractor Reconstruction from Time-Series Data	69
Fractals and Combat	70
Evidence of Chaos in War From Historical Data?	71
Evidence of Self-Organized Criticality From Historical Data?	73
Use of Complex Systems Inspired Measures to Describe Combat	75
Use of Relativistic Information to Describe Command and Control Processes	78
Tier V: Combat Technology Enhancement	80
Computer viruses ("computer counter-measures")	80
Fractal Image Compression	80
Cryptography	83
Tier VI: Combat Aids	85
Using Genetic Algorithms to Evolve Tank Strategies	86
Tactical Decision Aids	89
Tactical Picture Agents	94
Autonomous Robotic Devices	96
Tier VII: Synthetic Combat Environments	97
Combat simulation using cellular automata	97
Agent-based simulations	98
Swarm	102
Tier VIII: Original Conceptualizations of Combat	105
Dueling Parasites	105
Percolation Theory and Command and Control Processes	107
Exploiting Chaos	108
Pattern Recognition	111
Fire-Ant Warfare	117
Suggested Directions	119
General Directions	119
Some Open Questions	122

Appendix A: A Summary of the Mathematical Tools of Nonlinear Dynamics and Complex Systems Theory Discussed in Part I of this Report	127
Appendix B: Nonlinear Dynamics and Chaos	131
Key Concepts	131
Basic Lessons	133
Appendix C: Complex Systems Theory	137
Key Concepts	137
Basic Lessons	139
Appendix D: Irreducible Semi-autonomous Adaptive Combat	
Agents	143
What is the basic idea?	144
Actions, Goals, Properties, and States	145
Decision Dynamics	150
Measures of Combat effectiveness	151
Appendix E: A Brief Discussion of the General Applicability of Genetic Algorithms	153
References	157

Executive Summary

The Commanding General, Marine Corps Combat Development Command (MCCDC), asked the Center for Naval Analyses (CNA) to assess the general applicability of the *new sciences* to land warfare. "New sciences" is a catch-all phrase that refers to the tools and methodologies used in nonlinear dynamics and complex systems theory to study physical dynamical systems exhibiting a "complicated dynamics."

A Concise Summary of the Overall Assessment

This report concludes that the concepts, ideas, theories, tools and general methodologies of nonlinear dynamics and complex systems theory show enormous, almost unlimited, potential for not just providing better solutions for certain existing problems of land combat, but for *fundamentally altering our general understanding of the basic processes of war, at all levels*. Indeed, the new sciences' greatest legacy may, in the end, prove to be not just a set of creative answers to old questions but an *entirely new set of questions* to be asked of what really happens on the battlefield.

This most far-reaching assessment necessarily represents a long-term view of the coupling between the new sciences and military theory. It assumes that a reasonably long-term commitment will be made to ensure that at least some of the many possible evolutionary routes borne of this coupling can be adequately explored. New ideas, particularly those that question the very foundation of a discipline, require time to mature. Therefore, one should not necessarily expect to see a "killer-application" that completely revolutionizes the way we fight wars any time soon.

At the same time, however, and from a shorter-term perspective, this report concludes that many of today's "conventional" problems and issues of warfare – such as the dissemination and fusion of real-time battlefield data; tactics and/or strategy development; and more intelligent use of models and simulations – can benefit, to varying degrees, from the basic insights into the behavior of complex systems obtained by the new sciences.

Central Thesis

The central thesis of this paper is that *land combat is a complex adaptive system*. That is to say, that land combat is essentially a nonlinear dynamical system composed of many interacting semi-autonomous and hierarchically organized agents continuously adapting to a changing environment.

Table 1. Land combat as a complex adaptive system

Generic Property of Complex Systems	Description of Relevance to Land Combat
<i>Nonlinear interaction</i>	Combat forces composed of a large number of nonlinearly interacting parts; sources include feedback loops in C2 hierarchy, interpretation of (and adaptation to), enemy actions, decision making process and elements of chance
<i>Nonreductionist</i>	The overall "fighting ability" of a combat force is not a simple aggregate function of the fighting ability of individual combatants
<i>Hierarchical structure</i>	Combat forces organized in a command and control hierarchy
<i>Decentralized control</i>	There is no master "oracle" dictating the actions of each and every combatant
<i>Selforganization</i>	Local action, which often appears "chaotic" induces long-range order
<i>Nonequilibrium order</i>	Military conflicts, by their nature, proceed far from equilibrium
<i>Adaptation</i>	In order to survive, combat forces must continually adapt to a changing environment
<i>Collectivist dynamics</i>	There is a continual feedback between the behavior of (low-level) combatants and the (high-level) command structure

Military conflicts, particularly land combat, have almost all of the key features of complex adaptive systems (see table 1): combat forces are composed of large numbers of nonlinearly interacting parts and are organized in a command and control hierarchy; local action, which often appears disordered, induces long-range order (i.e. combat is self-organized); military conflicts, by their nature, proceed far from equilibrium; military forces, in order to survive, must continually adapt to a changing combat environment; there is no master "voice" that dictates the actions of each and every combatant (i.e. battlefield action effectively proceeds according to a decentralized control); and so on. This means, in principle, that land combat ought to be amenable to

precisely the same methodological course of study as any other complex adaptive system, such as the stock-market, a natural ecology or the human brain.

Implicit in this central thesis is the idea that these largely conceptual links between properties of land warfare and properties of complex systems in general can be extended to forge a set of practical connections as well. That is to say, land warfare does not just *look like* a complex system on paper, but can be well characterized in practice using the same basic principles that are used for discovering and identifying behaviors in complex systems.

Eight Tiers of Applicability

The following eight tiers of applicability provide a convenient scaffolding on which to organize the potential applications of complex systems theory to warfare. Indeed, just the fact alone that there are so many different levels to which the "new sciences" can be applied testifies to their enormous potential. Note that *Tiers I* through *VIII* range roughly from applications involving the least risk and least potential payoff (at least, as far as a practical applicability is concerned) on *Tier-I*, to applications involving the greatest risk, but also the greatest potential payoff, on *Tier-VIII* (see figure 5):

Tier-I: *General metaphors for complexity in war*

The first tier of applicability consists of constructing and elaborating upon similar sounding words and images that most strongly suggest a "philosophical resonance" between behaviors of complex systems and certain aspects of what happens on a battlefield. It is on this tier that the well-known Clauswitzian images of "fog of war," "center-of-gravity" and "friction" are supplanted by such metaphors as "nonlinear," "co-evolutionary" and "emergent." This first tier is accompanied by words of both encouragement and caution:

- On the one hand, the act of developing metaphors is arguably an integral part of what complex systems theory itself is all about, and therefore ought to be encouraged
- On the other hand, an unbridled, impassioned use of metaphor *alone*, without taking the time to work out the

details of whatever deeper insights the metaphor might be pointing to (i.e. without exploring what the other tiers of applicability might have to offer), runs the risk of both shallowness and loss of objectivity.

Tier-II: Policy and General Guidelines for Strategy

The second tier of applications takes a step beyond the basic metaphor level of *Tier I* by using the metaphors and basic lessons learned from complex systems theory to guide and shape how we formulate strategy and general policy. *Tier-II* thus extends the first tier of application to the military organization as a whole. It consists of using both the imagery of metaphors and the tools and lessons learned from complex systems theory to enhance and/or alter organizational and command and control structures. Potentially useful policy implications of ideas borrowed from the lessons of complexity theory include

- *Look for Global Patterns.* Search for global patterns in time and/or space scales higher than those on which the dynamics is defined. Systems can appear to be locally disordered but still harbor a global order.
- *Exploit Decentralized Control.* Encourage decentralized control, even if each "patch" attempts to optimize for its own selfish benefit, but maintain interaction among all patches.
- *Find Ways to Adapt Better.* The most successful complex systems do not just continually adapt, they struggle to find ways to continue to adapt better. Move towards a direction that gives you more options.

Tier-III: Conventional warfare models and approaches

Tier-III consists of applying the tools and methods of nonlinear dynamics to more or less "conventional models" of combat. The idea on this tier is not so much to develop entirely new formulations of combat so much as to extend and generalize existing forms using a new mathematical arsenal of tools. Examples of applications include

- Using nonlinear dynamics to explore implications of nonlinearities in generalized forms of the Lanchester equations
- Exploiting an analogy between the form of the Lanchester equations and Lottka-Volterra equations describing predator-prey interactions in natural ecologies to develop new models of combat
- Using genetic algorithms to perform sensitivity analyses and otherwise "test" the veracity of existing complex simulation models

Tier-IV: *Description of the complexity of combat*

This tier consists of using the tools and methodologies of complex systems theory to describe and help look for patterns of real-world combat. The fundamental problem is to find ways to identify, describe and exploit the latent patterns in behavior that appears, on the surface, to be irregular and driven by chance. Examples of applications include

- Looking for evidence of chaos in historical combat data
- Using various qualitative and quantitative measures from nonlinear dynamics and complex systems theory to describe the complexity of combat
- Use phase-space reconstruction techniques from nonlinear dynamics to reconstruct attractors from real-world combat data and make short-term predictions based on underlying patterns

Tier-V: *Combat technology enhancement*

Tier V consists of applying complex systems theory tools to enhance existing combat technologies. The objective of this middle tier of applications is to find ways to improve, or provide better methods for applying, specific key technologies. Examples of applications include

- Using fractals for data compression

- Using cellular automata and chaotic dynamical systems for cryptography
- Using genetic algorithms for intelligent manufacturing
- Using synchronized chaotic circuits to develop cheap IFF

Tier-VI: Combat aids for the battlefield

Tier-VI consists of using the tools of nonlinear dynamics and complex systems theory to enhance real-world operations. Examples of applications include

- Using genetic algorithms to "evolve" operational tactics and targeting strategies
- Developing tactical picture agents to adaptively identify, filter and integrate relevant information in real-time
- Developing autonomous robotic devices to act as sentries and to help in material transportation and hazardous material handling

Tier-VII: Synthetic combat environments

Tier-VII consists of developing full system models for training purposes and/or for use as research laboratories from which general (and possibly universal) patterns of behavior can be obtained. Examples of applications, ranging from least to most sophisticated, include

- Using cellular automata to explore basic behavioral properties of simple local rule-based combat models
- Using multi-agent based simulations of combat to explore behavioral properties of combat models of mid-level complexity
- Using the *Santa Fe Institute's* general purpose modeling system called SWARM to develop a full system-level model of land warfare

Tier-VIII: *Original conceptualizations of combat*

Tier VIII represents the potentially most exciting – and certainly most far-reaching – tier of the eight tiers of application. It consists of using complex systems theory inspired ideas and basic research to develop *fundamentally new conceptualizations of combat*. Examples of applications include

- Using genetic algorithms to "evolve" possible low-level rules that describe high-level observed combat behavior
- Using neural nets to "induct" otherwise unseen behavioral "patterns" on a battlefield
- Developing ways of exploiting the nature of chaos in combat phase-space to selectively "drive" combat to move towards more favorable regions
- Exploiting the collective intelligence of very many otherwise "simple" autonomous micro-bots to conduct "Fire-Ant Warfare"

Most Promising Applications

In a roughly ascending order of the probable length of time that an application is likely to require before maturing to a point at which a definitive assessment of its payoff can be made, here are seven of the most promising applications of the "new sciences" to land warfare that can be made in the short-term:

1. Exploit the general analogy between the form of the Lanchester equations and Lotka-Volterra equations describing predator-prey interactions in natural ecologies to develop a generalized "Neo-Lanchesterian" approach to land combat.
2. Use phase-space reconstruction techniques from nonlinear dynamics to reconstruct attractors from real-world combat data and make short-term predictions based on underlying patterns. Part of this involves developing an appropriate phase-space description of combat.

3. Develop simple local-rule-based models of combat to explore general behavioral patterns and possible universalities in combat.
4. Use genetic algorithms to "evolve" tactics and strategies. Specific applications might include tank tactics, targeting strategies and using genetic algorithms as backbones of real-time adaptive battlefield decision aids.
5. Develop multi-agent-based simulations of land warfare to be used as training tools along the lines of commercial games such as *SimCity* and *SimLife*. Explore the possibility of using the Santa Fe Institute's SWARM modeling system.
6. Develop agent-based tactical picture agents to adaptively retrieve, filter, and integrate battlefield and intelligence data.
7. Reexamine existing policy and policy procedures, *at the highest levels*, in light of the basic lessons learned from complex systems theory.

Of course, there are many more theoretical avenues to explore in the long-term as well. These include developing measures of "complexity" of combat, developing general data-collection methods that emphasize "process" vice more traditional force-on-force attrition "statistics," looking for and exploiting characteristic fractal-like behaviors in combat, using various sophisticated pattern recognition techniques to look for any high-level exploitable patterns on the battlefield and/or information databases that describe the progress of a campaign, and finding ways of exploiting the ability to both "control" and "tame" chaos on the battlefield. These, and other possibilities, are all discussed in the main text of this paper.

As a final note, it must be emphasized that land combat represents but one level of activity within a complex nested hierarchy of levels existing on many scales, not the least of which is political. To make full use of what complex systems theory has to say about the general nature of warfare, its lessons must be applied not just to land combat alone, but to the entire chain of combat and command structures.

General Guidelines

The most important overall suggestion that can be made regarding the applicability of complex systems theory to land warfare is to *be patient!* As discussed at length in Part I of this report, and stressed repeatedly throughout both volumes, complex systems theory is a very young, very immature science, which – at this time – is not even sure of its *own* future direction, much less of its applicability to other, specific areas.

General guidelines for applying some of the basic lessons learned from nonlinear dynamics and complex system theory include:

- **Develop "Nonlinear Intuition."** It is vital for every decision maker to go beyond the conventional "linear" intuition and develop an intuition for the kinds of nonlinear behaviors pervasive in complex systems. If the ideas of nonlinear dynamics and complex systems theory are to take root in the military, it is important for its leaders to learn some of the technical aspects of these approaches.
- **Look for Inherent Nonlinearities in Conventional Models.** A fundamental lesson of nonlinear dynamics theory is that one can almost always expect to find some manifestation of chaos whenever nonlinearities are present in the underlying dynamics of a model. This fundamental lesson has potentially significant implications for even the simplest combat models. Though some work has recently been done to determine the implications of having nonlinearities embedded within conventional models, many important insights into how our current models of land combat really behave remain to be discovered.
- **Emphasize Strong Interdisciplinarity.** If there is one universally agreed upon "insight" that has emerged out of *Santa Fe Institute's* first 12 years of existence it is that progress in complex systems theory demands an interdisciplinary approach. Complex adaptive systems are best studied by other complex adaptive systems.
- **Redefine Traditional Measures of Effectiveness and Data Collection Requirements.** If land combat is a bona fide candidate system for study as a complex system it must, initially, be treated essentially as "just another system" for

study by complex systems theory. This means that such basic questions as *"What are appropriate MOEs?"* and *"What kinds of data are required for understanding the processes of combat?"* must be re-examined from the point of view of complex systems theory.

- **Do not Shy Away from using "Simple" Models.** The first task of any fundamental research effort – and this is what finding ways of applying complex systems theory to land combat must necessarily be viewed as at this juncture – is to find a simple enough system that, while it is not an exact replica of the system that one is trying to understand and may lack many of its real-world complications, is able to capture some of the essential properties of the real system.
- **Attack Problem from Diverse Fronts.** Complex system theory, in practice, consists of looking for all possible ways to gain a better understand of how a complex system behaves. This means attacking the problem from diverse fronts: develop models of land combat, compare behaviors of combat models with behaviors of models of other systems, and develop new tools to record relevant data (and to re-examine historical data) from a complex systems theory perspective.

General Discussion

Preliminary Remarks

In the opening section of *Land Warfare and Complexity, Part I* [28], a not entirely facetious question was raised asking whether the behavior of the human brain, with its enormously complicated set of about ten billion nonlinearly interacting neurons, has anything in common with what happens on the battlefield? On paper, the human brain and the battlefield appear to have much in common. For example, both consist of a large number of nonlinearly interacting parts whose individual behaviors depend on the action and pattern of behavior of other (nearby and not-so-nearby) parts. Both obey a decentralized control. Both appear to be locally "chaotic" but harbor intricate, long-range order. Both tend not to dwell for long times near equilibrium, preferring instead to exist almost exclusively in a nonequilibrium state far from equilibrium. Both continually adapt to internal and external pressures and to the environment. And so on. But the fact there are obvious core similarities in the makeup and dynamics of these two a-priori very different systems is not really the most interesting question one can ask. The really interesting question is "*What universal patterns of behavior do these two a-priori dissimilar systems produce by virtue of their core similarities?*" And are there any universal patterns of behavior common to all complex adaptive systems that share certain "signature" characteristics; and if so, what are those characteristics?

Scott Kelso, who is the director of the *Center for Complex Systems* at Florida Atlantic University, has recently published an ambitious and intellectually far-reaching monograph exploring the thesis that the human brain is fundamentally a pattern-forming, self-organized dynamical system obeying a set of nonlinear dynamical laws [32]. In this monograph, he writes,

"Rather than compute, our brain 'dwells' (at least for short times) in metastable states: it is poised on the brink of instability where it can switch flexibly and quickly. By living near criticality, the brain is able to anticipate the future, not simply react to the present. All this involves the new physics of self-organization in which ... no single level is any more or less important than any other."

It is tempting to speculate that almost exactly the same words can be ascribed to what happens on a battlefield.

Another parallel can be drawn between Kelso's book and Parts I and II of this assessment. Just as Kelso's book is unarguably an *anomaly* from the standpoint of the conventional practices of cognitive- and neuro-sciences – sciences that have heretofore used tools and methodologies rather far-removed from the realm of nonlinear dynamics and complex systems theory – so, too, is a discussion of applying the "new sciences" to the study of land warfare an anomaly if interpreted in the context of conventional military theory, operations research, and modeling practices. (Not surprisingly, Kelso's book – and the overall approach it advocates – has thus far garnered only a lukewarm reception from the practitioners of the prevailing wisdom.) The lesson here is that conventional wisdom is usually a difficult and stubborn beast to nudge. The assertion that *land combat is a complex adaptive system*, along with the general methodology that this paper proposes be used to explore the implications of this assertion, is as radical a suggestion to make in operations research circles as Kelso's thesis is in the circle of cognitive- and neuro-scientists. The skeptical reader is therefore urged to give this novel and admittedly far-from-mainstream idea an honest chance.

The Charter

The titles of Part I [28] and Part II of this paper should serve as strong indicators of the fact that the charter for this project was very broad. This project was designed to *assess the applicability of nonlinear dynamics and complex systems theory to the study of land warfare*. The form of the overall assessment, and the level on which the overall assessment of applicability was to be made, was deliberately left open. There are two main reasons for this. One reason, of course, is the relative immaturity of the so-called "new sciences," in which both nonlinear dynamics and complex systems theory are major players. The second reason is that, as alluded to in the opening remarks to this section, this is the first serious attempt that any branch of the military has made to incorporate into its world-view such a broad spectrum of knowledge that heretofore has been largely dismissed as being either only of vague, peripheral interest or as being utterly irrelevant altogether. In light of these two reasons, expectations and constraints were both kept to a minimum.

The Problem

Nonlinear dynamics and complex systems theory have been around as bona-fide research disciplines for roughly 30 years and 12 years, respectively. Both are still maturing fields. Particularly complex systems theory, which cannot at this time be considered anything but an "infant science." Together, nonlinear dynamics and complex systems theory form a growing pool of knowledge and conjectures about (1) what tools are best for describing the characteristics of real-world complex systems and for describing real-world systems that exhibit an apparently "complicated" dynamics, and (2) what general behavioral properties many real-world complex systems all seem to share. Thus, the basic question that was asked at the beginning of this project was:


"Given what we thus far know about how to study real-world complex systems and how those systems generally behave, what do nonlinear dynamics and complex systems theory add to our conventional understanding of land warfare?"

On one level, the most honest and simplest answer one can give to this question is that "We'll know as soon as we start looking at land warfare in earnest." For reasons already discussed, complex systems theory is not yet a mature enough "science" that it is able to say much of anything about a system other than identifying it as a viable candidate for study. The fact that one application of complex systems theory to land warfare might involve, say, a multi-agent simulation, *by itself* tells us little about land warfare that we did not know already. Even what is meant by "multi-agent simulation" is not all that clear, with different researchers choosing to interpret their own favorite meanings. What potentially new insights into land warfare a multi-agent simulation might provide us with in the future requires us to first develop and then study in detail the behavior of such a model. *It is, in fact, absolutely vital that we take such a step*, for in the end there is no substitute for seeing what actually comes of interacting with a complex systems theory model of land combat. If recent history of complex systems theory-derived models of natural systems gives any indication as to what kinds of insights to expect from such models, the best guess is that the most important insights will prove to be unexpected and emergent. Recall Conway's simple two-dimensional cellular automaton rule called *Life* (Part I, page 87). While the rule itself is very simple,

and is easy to set up to run on a computer, it took years for researchers to catalog its many patterns and to eventually prove that it is capable of universal computation.

On a deeper level, the answer to the boxed question on the previous page really embodies three separate but interrelated issues (see figure 1):

Figure 1. Schematic of the interrelated issues of addressing land warfare as a complex system



1. *Complexity theory*, which refers to any and all conjectures, hypotheses, theories, experiments, mathematical models, etc. having to do with the understanding of complex systems exhibiting a complicated (i.e. chaotic) behavior. In particular, complexity theory is assumed to include both nonlinear dynamics and complex systems theory, the latter including a multitude of sub-disciplines such as artificial life, cellular automata, genetic programming, neural networks, etc. Part I [28] of this report contains a detailed discussion of many of these topics. Topical summaries are also provided both in the discussion below and in the appendix.
2. *Land warfare*, which embodies all of the myriad problems and issues of land warfare, including combat attrition, command and control, coordination, intelligence, tactics and strategy, training, etc. One must also respect that the various levels of land warfare to which the tools and methodologies of complexity theory can be applied, including *tactical, operational, strategic* and *general strategic* levels.¹

¹ The general strategic level refers to the level on which socio-political strategies are followed over long periods of time and which can therefore span over several conflicts and embody many different strategies.

3. *Modeling/Simulation*, which is a generic label for the overarching context within which possible interconnections between the tools and methodologies of complexity theory as well as the issues and problems of land warfare can be fully explored. One must be cognizant of the overall objective of any model or simulation before it is developed. A model that is to be used only for training purposes is fundamentally different from a model that is to be used to assess the effects of weapons and/or decision-making in a given environment.

The Approach

Before an honest effort could be made to discuss and weigh the relative merits of the possible applications of the so-called "new sciences," several preliminary information-gathering steps had to first be taken:

1. A careful and extensive review of all available technical and popular literature
2. A review of information resources available on the World-Wide-Web (WWW)
3. A review of the state-of-the-art technology and methodology, as practiced by active researchers in the fields of nonlinear dynamics and complex systems theory
4. Consultation with research staff members of the *Santa Fe Institute* in Santa Fe, New Mexico

Part I of this report provides the theoretical framework and mathematical background necessary to understand and intelligently discuss many of the key ideas and concepts underlying the study of nonlinear dynamics and complex systems theory. It is also meant to be consulted as a general technical sourcebook of information. Part I includes an extensive bibliography and glossary of terms, and a sorted collection of many nonlinear dynamics and complex systems theory related Universal Resource Locator (URL) links on the WWW. In fact, during the write-up of Part I, it was decided that an HTML-formatted version of the glossary of terms and a somewhat shortened version of the URL list of WWW resources could serve

as a cornerstone of MCCDC's new "New Sciences" home page on the WWW.²

Two Preliminary Questions

From the title of this paper, and from the basic charter of this project, one can intuit that before any prolonged discussion of potential applications ensues, two basic questions must first be answered:

1. *What are nonlinear dynamics and complex systems theory?*
2. *Why land warfare and not air warfare or naval warfare?*

Question 1: What are Nonlinear Dynamics and Complex Systems Theory?

In simplest possible terms, "nonlinear dynamics" refers to the study of dynamical systems that evolve in time according to a nonlinear rule. This means that, for example, the effect of adding two inputs first and then operating on their sum is, in general, not equivalent to operating on two inputs separately and then adding the outputs together. Or, more colloquially, the whole is not necessarily equal to the sum of the parts. Complex systems theory, on the other hand, refers to the study of dynamical systems that are composed of many nonlinearly interacting parts.

Nonlinear dynamics and complex systems theory both fall under the broad rubric of *complexity theory* that embodies a remarkably wide variety of disciplines ranging from biology, chemistry, and physics to anthropology to sociology to economics. Among the many subfields of complexity are deterministic chaos, stochastic dynamics, artificial life, ecological and natural evolutionary dynamics, evolutionary and genetic programming, cellular automata, percolation theory, cellular games, agent-based modeling, and neural networks, among many others (see table 2). Many of these have been either discussed in length or touched upon briefly in Part I of this report [28].

² MCCDC's "New Sciences" homepage can be found at this address:
<http://138.156.204.100/www/MCRC/library/beyond.htm>

Complexity theory as applied to land combat is best thought of as a process whereby one searches for the simplest possible description of combat that gives rise to the greatest variety of real-world behaviors. The supposition is that there are *universal* behaviors that do not depend critically on the details of force structure and dynamics.

Table 2. A small sampling of research areas, concepts and tools falling under the broad rubric of "complexity"

Research Areas	Concepts	Tools
agent-based simulations	adaptation	agent-based simulations
artificial life	autonomous agents	backpropagation
catastrophe theory	autopoiesis	cellular automata
cellular automata	complexity	cellular games
cellular games	computational irreducibility	chaotic control
chaos	computational universality	entropy
chaotic control theory	criticality	evolutionary programming
complex adaptive systems	dissipative structures	fuzzy logic
coupled-map lattices	edge-of-chaos	genetic algorithms
discrete dynamical systems	emergence	inductive learning
evolutionary programming	fractals	information theory
genetic algorithms	intermittency	Kolmogorov entropy
lattice-gas models	phase space	lattice-gas models
neural networks	phase transitions	Lyapunov exponents
nonlinear dynamical systems	prisoner's dilemma	maximum entropy
percolation theory	punctuated equilibrium	neural networks
petri nets	self-organization	Poincare maps
relativistic information theory	self-organized criticality	power spectrum
self-organized criticality	strange attractors	symbolic dynamics
time-series analysis	synergetics	time-series analysis
<i>etc.</i>	<i>etc.</i>	<i>etc.</i>

Despite the fact that there is considerable overlap both between nonlinear dynamics and complex systems theory, and among the individual research areas, concepts and tools that constitute these two overlapping disciplines, there are two deep themes that run through, and summarize the essence of, *all* complexity research:

- *Surface complexity can emerge out of a deep simplicity*, embodying the idea that what may at first appear to be a complex behavior, or set of behaviors, can in fact stem from a simple underlying dynamics
- *Surface simplicity can emerge out of a deep complexity*, embodying the idea that enormously complicated systems that *a-priori* have very many degrees-of-freedom and

therefore are expected to display "complicated" behavior, can, either of their own accord via self-organization, or through selective "tuning" by a set of external control parameters, behave as though they were really low-dimensional systems exhibiting very "simple" behavior

Most "complexity research" consists of attacking specific problems from both of these points of view simultaneously. This is also the method of attack we will use throughout this paper in applying complex systems theory to the problem of land warfare.

Question 2: Why Land Warfare and not Air Warfare or Naval Warfare?

Of the many different kinds of modern warfare that could be chosen as a testbed for applying the ideas and tools of complex systems theory, land warfare – *as a whole* – is best suited for three reasons:

1. **Lots of Individual Components.** Land warfare involves potentially vast numbers of mutually interacting combatants (where "combatant" refers to any irreducible element of combat; i.e. it can refer to an infantryman, a tank or transport vehicle, depending on circumstance), so that it can naturally be described as a "complex system." While a natural unit of measure for the size of a naval force, for example, is the number of ships, the natural unit for land warfare is the number of *individual soldiers*.
2. **Complex Environment.** Land warfare generally takes place within a much more complex environment than do other forms of combat. Air combat, for example, takes place within an almost homogeneous medium and "interactions" arise mainly from line-of-sight. Likewise, while the ocean is arguably a complicated medium, from the point of view of vulnerability to enemy fire, it represents a relatively simple environment. On the other hand, the surface of the earth is strewn with complexities, from its effects on various sensors and communications systems to its profound implications for the composition of combat forces and tactics.

3. **Involves Individual Psychology.** Land warfare depends on psychological factors to a far greater extent than do other forms of combat.

Now, these three reasons ought not be misunderstood to imply that aspects of other kinds of warfare are not amenable to study from a "new sciences" perspective. Such a view is, in fact, patently false. For example, questions regarding how to best understand the operation of command and control structures and vulnerabilities of enemy Integrated Air and Defense Systems (IADS), are inherently imbued with various complexity-laden problems that are best tackled by applying the tools and methodologies of complex systems theory. Likewise, the integration of many simultaneous real-time streams of information is an integral part of putting together a reliable and meaningfully evolving tactical picture. This is a fundamental problem that is amenable to some of the ideas borne of complex systems theory and its "solution" is applicable, in principle, to all forms of combat.


What Have We Learned From Complexity Theory?

Just as "chaos theory" is a misnomer because the theory really deals with the regularities and patterns embedded within what only appears to be disordered and chaotic, so "complexity theory" is a misnomer because it really deals with the underlying simplicity of what are ostensibly enormously complicated systems. In fact, complexity theory attacks the problem of describing the behavior of complex systems from two complementary points of view (see figure 2):

- *Surface complexity arising out of a deep simplicity*, in which it attempts to describe the apparently complex high-level behavior of a complex system (such as the language capability of the human brain) in terms of a much simplified set of low-level rules (such as the neural-net-like threshold activation functions of individual neurons of a human brain). Think of the two-dimensional cellular-automaton *Life*-game (see Part I, page 87) in which complex high-level patterns of gliders and other self-reproducing structures owe their existence to a rather simple set of underlying rules.

- *Surface simplicity arising out of a deep complexity* (Kelso, [32]), in which it tries to elucidate the general principles of how complex systems can be made to behave as though they really lived in a much simpler (or lower-dimensional) space. Think of the Rayleigh-Benard experiment [23] in which a layer of liquid (composed of a vast number of interacting molecules) is heated from below and cooled from the top to set up a temperature gradient. If this temperature gradient is small, heat is randomly dissipated among the many molecules and there is no large-scale patterned motion of the liquid. As the gradient increases, however, the liquid begins moving as a coordinated whole in an ordered, coherent, rolling fashion. Its overall behavior can now be expressed by a single collective variable describing the roll-amplitude. In other words, the behavior of a complex system that a-priori has on the order of 10^{20} degrees-of-freedom (since there are about 10^{20} molecules, each of which contributes to the overall motion), can be described by a one-dimensional equation of motion!

Figure 2. Two complementary lessons of complex systems theory



What We Know

Despite the infancy, as a science, of complex systems theory, and the relative dearth of universally applicable results, there is nonetheless much that has been learned about the general behavior of complex adaptive systems in recent years.

We know, for example, that in order for complex systems to be "adaptive" it is necessary that they are *nonlinear* and capable of both *storing* and *exchanging information* among their parts. The

information exchange must also be meaningful; that is, it must be neither too sparse nor too dense (recall the *edge-of-chaos* metaphor, discussed in Part I, page 76).

We also know, or at least have a very strong suspicion, that *self-organization* is a general property of complex adaptive systems. This means that complex adaptive systems are such that, as they evolve, they naturally go from being initially chaotic, featureless, and possessing disorganized, independent, states to having organized, structured and highly inter-dependent states. This progression from order to disorder may be smooth or proceed with many fits and starts. It may even reverse direction at times, as it does in natural evolution. But the overall tendency in any complex adaptive system is to evolve towards self-organization.

There are roughly seven basic concepts of and conditions for self-organization in complex systems ([32]):

1. The spontaneous appearance of patterns results from large numbers of *nonlinearly interacting components*. If the system does not possess a sufficient number of components, or if its components do not interact, patterns will not emerge.
2. The system must be *dissipative* (that is, there must be mechanisms for converting one form of energy into another, and of locally pumping energy out of the system) and *far from equilibrium*. Because of the nonlinear interactions, heat or energy does not diffuse uniformly throughout the system. Rather, energy is concentrated into structural flows that transport the heat (dissipate it) more efficiently. As a result of this dissipation, many of the systems degree-of-freedom are effectively suppressed and the system behaves as though it lives in a much lower dimensional space.
3. *Control parameters* lead the system through different patterns but are typically not dependent on the patterns themselves. In the Rayleigh-Benard experiment (see above), the control parameter is the temperature gradient that is set up between the top and bottom layers of the liquid.
4. *Collective variables* (or *order parameters*) describe the relevant degrees-of-freedom of the system. These collective

variables are created by the coordination among the system's parts. In the Rayleigh-Benard experiment (see above), the collective variable is the amplitude of the convection rolls of the liquid.

5. Order parameters are usually found near *nonequilibrium phase transitions*, where a loss of stability entails new and/or different patterns and/or switching between patterns.
6. *Fluctuations* – which are an inherent source of internal noise – continuously probe the system, allowing it to adjust its behavior and search for new patterns.
7. The *order parameter dynamics*, or equations of motion describing the coordinated dynamical motion of the system, can have simple (i.e. fixed point or limit cycle) or complicated solutions, including deterministic chaos.

What We Don't Know

There is much about the behavior of complex adaptive systems that we do not know. We do not even know the answers to some of the most basic questions one could pose about complex systems, such as "What is complexity?" or "What is organization?" We do not know, for example, why some systems are able to adapt well to certain environments and/or environmental changes and others are not; i.e. we do not really know what makes a system adaptive. We do not know whether there is a minimum level of complexity a system must possess in order to be adaptive. or how to define a universally relevant measure of complexity. We do not know in advance whether a system will be weakly or strongly adaptive. We do not know exactly what properties of systems are characteristic of universal behaviors, nor what that class of universal behaviors is (*self-organized criticality* – see Part I [28], page 101 – notwithstanding).

How is Work in the "New Sciences" Actually Done?

We briefly discuss how the "new sciences" are really done in the complex systems theory research community to give a feel for what a complex systems theoretic approach to studying land combat is likely to look like.

There is a growing popular misconception that complex systems theory is a well-defined science. Somewhat facetiously, it could be said that a prevailing belief is that complex systems theory consists of some canned set of software routines that can be downloaded from, say, *Microsoft's* WWW site, and directly unleashed on whatever "complex problem" happens to strike one's fancy. This cannot be further from the truth. The reality is that much of what goes under the name of "complex systems theory" actually consists of a hodgepodge of on-the-fly hand-crafted and tinkered techniques and approaches that say more about the research style of a particular complex systems "theorist" than they do about the how the new sciences are practiced as a whole. There is certainly no existing complex systems theory model *per se* that can be ported over to describe land combat. The current crop of models are either specifically tailored to particular problems – such as John Holland's ECHO is for ecological studies or Menczer's and Belew's Latent Energy Environment (LEE) is for investigating the general question of how the behavior of organisms is interconnected with their environment – or are general purpose simulators (like the *Santa Fe Institute's* SWARM) that must be carefully tuned to apply to specific systems.

Most "new sciences" research is practiced by following these five basic steps (these steps are not meant to be taken tongue-in-cheek!):

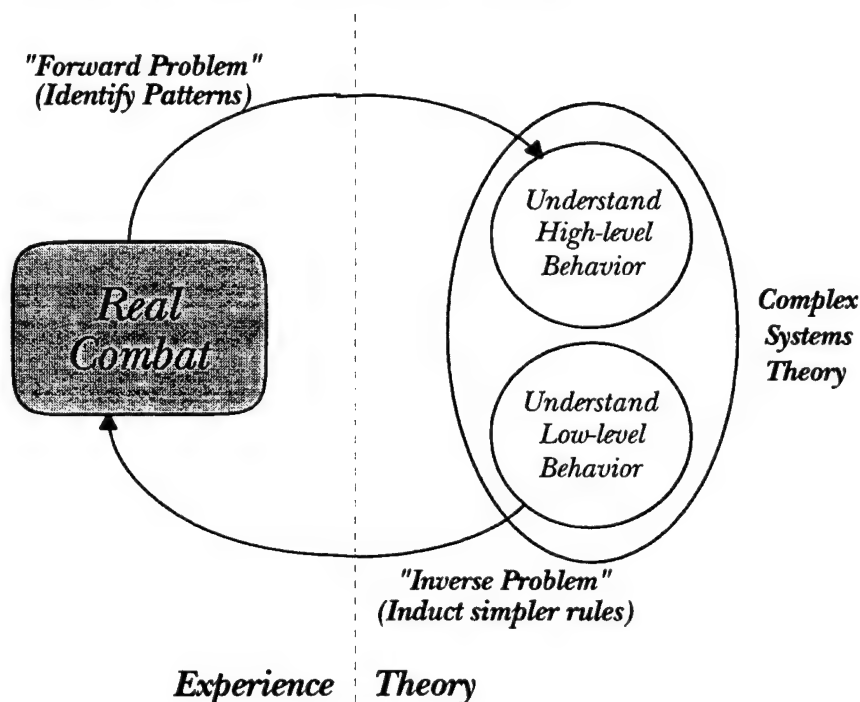
1. Think of an interesting question to ask regarding the behavior of a real system (or find a real system to study)
2. "Play" (i.e. interact) with simplified models of the system
3. *Sit back and watch for patterns*
4. Develop theories about how the real system behaves
5. Repeat steps 2-4!

The most important step is the italicized step 3. Much of the early work with trying to understand the behavior of a system consists of finding ways to spot overall trends and patterns in the behavior of a system while continually interacting and "playing" with "toy-models" of the system. If one is serious about applying the "new sciences" to land warfare, one must be ready to rethink

some of the conventional strategies and approaches to modeling systems. Complex systems theory is not necessarily best done by studying the output of a several-thousand-line long computer program.

Another important element of the basic approach of complex system theory to understanding the behavior of complex systems is that the "forward-problem" and "inverse-problem" must both be studied simultaneously (see figure 3), and that the interplay between experience and theory is never overlooked.

Figure 3. Interplay between experience and theory in the forward- and inverse-problems of complex systems theory



The "forward problem" consists essentially of observing either real-world behavior or the behaviors of a model of a complex system with the objective being to identify any emergent high-level behavioral patterns that the system might possess. The "inverse problem" deals with trying to induct a set of low-level rules that describe observed high-level behaviors. Starting with observed data, the goal here is to find something interesting to say about the properties of the source of the data. The forward problem is therefore concerned with theoretical tools that are used to *identify patterns*, while the inverse problem is concerned with tools that are used to *induct low-level rules* (or models) that generate the observed high-level behaviors.

What are the Basic Questions to ask of a Complex System?

Just as the first steps to take in understanding a system is to playfully interact with toy-model representations of it (see above), so the first few questions that a complex systems theorist typically asks of a new system under study are relatively simple (simple to ask, not necessarily to answer) and to the point:

1. What is the number and nature of the elementary constituents of the system?
2. What is the relative contribution each constituent makes to the overall integrity of the system (i.e. what is its "context" and a measure of its "fitness")?
3. What kind of internal and external sensory apparatus do the constituents have?
4. Do the constituents have a "memory"? How reliable is it, and how far back in time do they retain memories?
5. What is the nature of the interactions among the constituents?
6. How do the constituents communicate (i.e. how do they pass and process information)?
7. What can the constituents do? That is, what is their "action space"?
8. How are the constituents able to adapt to (both internal and external) changes in their environment?
9. If the system is externally controlled, what are the external control parameters?
10. If the system seems to exhibit a globally ordered (or coherent) pattern of behavior, what are the appropriate collective variables that describe that behavior?

These questions, and others like them, are likely to be among the first basic questions asked by any serious student of complex systems theory encountering "land combat as a complex system" for the first time.

Patterns on a Battlefield

"Themistocles said that a man's discourse was like to a rich Persian carpet, the beautiful figures and patterns of which can be shown only by spreading and extending it out; when it is contracted and folded up, they are obscured and lost." – Plutarch (46-120 A. D.)

At the most abstract conceptual level, the irrefutably complex processes of land combat still form an inherent *pattern*. And a pattern is something that, once understood or unfolded, can be exploited.

Implicit in any application of the "new sciences" to land combat – whether that application uses attractor reconstruction techniques to make short-term predictions, genetic algorithms to evolve operational tactics, or multi-agent based simulations – is the idea that there is some latent order underlying what appears on the surface to be irregular and chaotic. Indeed, science itself arguably takes its greatest steps forward whenever someone discovers a new pattern, or underlying unity, in the world where none was observed before.

Patterns are best discovered when we take a step back from something to focus on the thematic relationships among, and overall context for, the parts. The ability to *perceive* patterns in an otherwise patternless process is what we call *intuition*. For example, what lies behind the success of every successful wall-street broker is the broker's intuitive recognition of patterns underlying the succession of Dow-Jones Industrial stock prices (along with a recognition of patterns of correlations with other pertinent variables such as acquisitions and mergers). Likewise, what lies behind the success of every great field commander possessed of the mythic "battlefield intuition," is an intuitive recognition of patterns underlying the unfolding processes of combat on a battlefield.

A receiver in football, running into a strong head-wind with outstretched arms to catch a throw coming over his right shoulder forty yards from where the quarterback released the ball, does not need to know the exact starting coordinates and velocity of the ball and integrate Newton's equations in his head. He simply "knows" when to catch the ball! He can intuit – and exploit – a simple pattern, without working out the details of how that pattern came to be. Moreover, this intuition is borne of direct experiential data. The receiver's brain slowly, over the

course of many years of playing football, builds an internal "model" of the flight path of a football in varying situations and contexts.

Now, consider what happens to an inflated balloon as its untied end is released.³ It ricochets around the room in a frenzied chaos. Its path is extremely sensitive to initial conditions, and is thus prototypically chaotic. The equations governing its motion are nonlinear, so that unless the initial conditions are known exactly, they cannot be used in practice to predict the balloon's path very far into the future. How would a football receiver do in trying to catch it? Assuming the football player possesses an above-average degree of eye-hand coordination, it is likely that his brain would gradually – over many trials – induct a model of the balloon's meanderings. While his model will be far from perfect, of course, the football player will likely, over time, be able to catch the balloon say, 5 or 10 percent more frequently than what we would guess from chance alone. He will have intuitively learned to exploit the a-priori chaotic system's behavior *well enough* to consistently beat mere odds.

Having such an "learned" intuition of latent patterns would clearly pay great dividends on the stock market and is the reason why wall-street is now hiring so many physics and mathematics Ph.D.'s straight out of graduate schools. Doyne Farmer, one of the founding fathers, and most promising young stars, of nonlinear dynamics recently quit active research in the field to found the *Prediction Company*. The *Prediction Company* consists mostly of young Ph.D.'s and graduate students and is designed specifically for applying the ideas of nonlinear dynamics and complex systems theory to identifying exploitable patterns on the stock market. Though its algorithms are proprietary and the details of how well Farmer's company is doing have not been disclosed, Farmer has been quoted in recent interviews as suggesting that the *Prediction Company* is doing quite well as far as "keeping ahead of the game" is concerned. "Keeping ahead of the game" in the stock market, of course, might mean doing only a few percent better than average; but that is all it takes.

Combat as Soccer?

Now, what does this all have to do with warfare? This discussion was meant to motivate the basic idea that if nonlinear dynamics and complex systems theory do nothing else, they *provide an*

³


This nice example comes from reference [31], page 422.

arsenal of powerful pattern recognition tools. As a visual metaphor, consider the image of a soccer field shown in figure 4.

Soccer is not a bad metaphor to use for combat. Land combat is arguably a much closer spiritual cousin to games such as soccer and football than it is to chess. While chess is a rather slow, stoic, progression of individual moves and countermoves, combat consists of more fluid-like dynamical evolutions of a coherent cacophony of players.

Figure 4 is taken from a figure appearing in a recent *New Scientist* article describing the work of sports scientist Keith Lyons and his colleagues at the University of Wales Institute in Cardiff, England [44]. Lyons is able to show that while, on one level, the progression of moves in a soccer match may appear to be random – with the 22 players on both sides ostensibly free to run and kick the ball anywhere they like on the field – on another level they harbor highly structured innate patterns. Though figure 4 shows only the simplest such pattern – namely, the tendency of a Dutch goal keeper playing in the 1994 World Cup to kick mainly to the left-hand-side of the playing field - it points to much deeper insights.

Figure 4. Schematic representation of latent patterns on a soccer field⁴



Lyons uses *notational analysis* – which is a technique for transferring information about a game from video onto computer-generated grids – to study in detail the movement and styles of play of individual players. He is particularly interested in the way that individual playing style disrupts an opponent's overall pattern of play. There is evidence to suggest that the most successful teams are those that actively intervene to create and exploit perturbations to the "flow" of a game. The best teams are the ones that "maximize their chances" and can "break the game's structural constraints." [44] Moreover, Lyons has

⁴ This figure is reproduced from [44].

observed that, over the course of an entire season, teams undergo transitions among three basic phases: (1) a period of initial equilibrium in which they settle down to a comfortable style of play, (2) a period of intermittent chaos, during which they adjust to a perturbation (such as being asked to incorporate a new tactic or accommodate the playing style of a new player), and (3) a period of new equilibrium, in which a team's playing style settles into one that may or may not be as successful as that during the first period. The objective is to identify the types of perturbations – and *when to apply them* – in order to drive a team through the second, chaotic, phase into a third phase that is more successful than the first.

All of this, of course, translates almost directly to the problems and issues faced by a field commander. The overarching question is "What are the underlying patterns governing the processes of combat?" Answering this question from the point of view of complex systems theory entails looking at combat from a rather unconventional perspective. Rather than focusing on force-on-force attrition statistics and other static measures of *what has happened*, one must instead focus on patterns of attrition as they evolve over time and other dynamic measures of *what is in the process of happening*. While these ideas are necessarily vague at this stage, they underscore an important element of the attitudinal shift that accompanies any re-examination of an old "problem" from a "new sciences" perspective.

Models and Simulations: An Heuristic Discussion

"Cosmic Nonsense: Is any model of the universe that claims to be final and exhaustive." – E. M. Robinson (1987)

It is safe to say that, in its current stage of development, complex systems theory consists almost entirely of some form of modeling, simulation or the design of simulation-engines, such as the *Santa Fe Institute's* SWARM system, that run them. Quite literally, much of leading-edge complex systems theory is practiced by (1) choosing a real-world system to "understand," (2) building a model of it, and (3) watching and interacting with the model as it runs on a computer (while looking for patterns that might provide some deeper insight into what the real-world system itself really does). Inevitably, then, any serious discussion of the applicability of complex systems theory to land warfare – or, for that matter, the applicability of complex systems theory to the study of *any* real complex system – must include a discussion

of what is meant by "model" and "simulation." It is not always easy to make, or maintain, a proper distinction between these two important concepts.

The simplest heuristic distinction that can be made between "models" and "simulations" is that models are generally represented in terms of *equations*, and simulations are generally represented by *computer programs* [56]. Though not always true, models try to induct an overall understanding of a system by simplifying it (simplifications typically coming at the expense of realism), whereas simulations try to include as much detail about a system as possible to reproduce its overall behavior in a specified situation (the realism often coming at the expense of simplicity). A slightly more refined view to take of models and simulations is to think of a simulation as a kind of "upper level," or more sophisticated, model.

What Makes a Model Useful?

In order to be useful, a model ought to have at least these three fundamental ingredients:

- 1. A Model Must be Developed in a Well-Defined Context.** A model must start out with the basic question, "*What is this model needed for?*" A model that is designed without having an explicit, well-defined question or set of questions to be answered is at best dangerous and at worst useless. The answers to different questions require different models. Is the "answer" to be *qualitative* or *quantitative*? *Descriptive* or *predictive*? *Broadly applicable* or heavily *context-dependent*? Is the purpose of the model to explicitly mimic reality?; is it to create a synthetic environment to simulate the experience of battle?; is it to teach decision making?; or, is it simply to assess the relative merit of various options? What kind of a model is developed, and how detailed it is, depends strongly on how these, and other, context-defining questions are answered. Modelers often do not think enough about what purpose the model is supposed to serve, and pay too much attention to building more and more unnecessary detail into the model.
- 2. A Model must "Respectfully" Simplify the Real System.** A model ought to skeletonize (i.e. reduce) a system down to its most important parts and drivers without

compromising the overall integrity of the system. If the model turns out to harbor effectively the same order of complexity as the system it models, one may have succeeded in "canning" the system for future study – so that the system can be studied even when the system physically is not present – but the task of simplifying and understanding the behavior of the system remains effectively untouched (see discussion of *Tierra* below).

3. **A Model must Provide a "Shortcut" Solution.** If the purpose of a model is to simulate the evolution of a system, the model must be able to faithfully reproduce the system's essential behavior in a time fast-enough to allow whatever decisions must be made based on the model's output to be made on time. If the purpose of the simulation is to provide a real-time decision aid, for example, the model must "predict" the outcome of (possibly many) sets of possible starting conditions in less time than it takes one set of real conditions to evolve.

Three Classes of Models

Roughgarden, *et. al.*, have found it convenient to partition the set of possible models into the following three basic (and slightly overlapping) classes:

- **Class 1:** *Minimal Idea Models*
- **Class 2:** *Minimal System Models*
- **Class 3:** *Synthetic System Models*

Minimal Idea Models

A *Minimal Idea Model* (MIM) explores some idea or concept without the inconvenience of specifying the details of the system, the environment in which the system evolves, or much of anything else. The assumption is that the phenomenon of interest is a computational entity whose properties are essentially the same across a wide range of possible universes.

An example of a MIM is a simple cellular automaton model (see Part I, page 81) of a complex system. Recall that, in its most elementary form, a cellular automaton provides a way of

exploring the implications of having a discrete space, discrete time, and a discrete local dynamics, and *no more*. It can be used as a basic template on top of which more realistic models can be built, but is itself useful primarily for looking for possible universal behaviors that appear in all complex systems obeying a local-rule dynamics.

Because MIM's leave out many of the details of a real-world system, their success ought to be measured less in terms of their "predictive value" and more in terms of their ability to "make a point," demonstrate the plausibility of a concept or simply communicate an idea.⁵

Most models of complex systems in the complex systems theory community currently reside in this class. It is also likely that most of the early impact of applying complex systems theory to land warfare will come from this class of models.

Minimum System Models

A *Minimal System Model* (MSM) is designed to explore the dynamics of some greatly simplified subset of features of the real system and/or environment. An MSM is essentially a MIM with some attention given to modeling the real-world environment. This class of models respects the details of the real system, but judiciously strips away unnecessary information. Of course, it may turn out in the end that the omitted information was crucial for understanding how the real system behaves. What to include and what to exclude is always the design choice of the modeler. But if the omissions are carefully and wisely chosen, and the simplified system retains the essential drivers of the real system, an MSM is a useful vehicle from which to abstract basic patterns of behavior.

This is the class of models that is likely to constitute the "bread-and-butter" class of complex systems theory derived models of land combat; but only *after* a preliminary round of seeing what MIMs have to offer to modeling land warfare has been completed.

Synthetic System Models

A *Synthetic System Model* (SSM) is an expansion of an MSM in which (ideally) all the assumptions about, and known properties of, a natural system are treated formally. It is a synthesis of

⁵ See Chapter 24, page 433 in reference [7] for a further discussion of this point.

detailed descriptions of all the component parts and processes of the system of interest.

An example of an SSM is *Santa Fe Institute's* SWARM, with which it is possible to develop a full system-level description of land combat. The price to be paid for developing an SSM, however, is that the behavior of the SSM can be just as difficult to understand as the behavior of the real system. Which brings us to the next very important question...

Computer Models

"Computers are useless. They can give you only answers." – Pablo Picasso

The Lanchester equations of land combat were borne of a rich tradition in the sciences to build abstract, simplified models of natural systems. Such models tended to be analytical in nature, often taking the form of differential equations. The emphasis was on simplicity: such models provided simple descriptions of real processes and were generally simple to solve. With the advent of the computer, of course, modeling became more concerned with incorporating a greater and greater level of detail about a physical system. In fact, computer models offer the following important advantages over traditional forms of modeling:⁶

- **Computer models can capture real-world complications better than mathematical models.** While computer models are all, at heart, algorithmic prescriptions for carrying out the steps of a formal model – so that the distinction between model and computer-model is not really clear-cut – formal mathematical models are generally able to capture only the gross aggregate characteristics of a system (number of constituents, average properties, and so on). Computer models, on the other hand, are more adept at capturing the nuances that describe real-world systems. There is no easy way, for example, to use a mathematical model to describe a feedback between a local piece of information and a global variable.
- **Mathematical models can typically be solved only in the limit of infinite-sized populations and/or infinite time.** Mathematical models thus generally provide simplified idealizations of behavior, while computer models are able

⁶ See chapter 24 of reference [7].

to deal with the complications of having finite sized systems evolving for finite times.

- **Mathematical models are generally poor at describing transient behavior.** One of the most important general lessons of complex systems theory is that complex adaptive systems prefer to live in far-from-equilibrium states. The least interesting systems are those that reach an equilibrium. It is notoriously difficult to model far-from-equilibrium systems with traditional mathematical modeling techniques. Computer models are absolutely vital for studying transient behavior.
- **Computer models provide a controlled environment.** Computer models provide a controlled environment in which to interactively study effects of changing initial conditions, control parameters, boundary conditions, and so on. Mathematical models are much more inflexible.

One must not be overzealous in incorporating ever finer detail into a model, however...

What Price Complexity?

Suppose a model successfully reproduces the dynamics and behavior of a complex system *exactly*, which is arguably the best one could hope to do using any model. For example, suppose, hypothetically, that someone builds the perfect representation of land warfare. Two questions should immediately come to mind:

1. *How do you know that the purportedly "perfect model" is really perfect?* How do you know, for example, that each and every behavioral nuance of the "perfect model" is a perfect reflection of (i.e. is in a one-to-one correspondence with) the behavioral nuances of the real system?
2. *What do you now do with the model?* How do you now make use of what the "perfect model" is telling you in order to answer whatever questions you developed the model to answer?

Both questions are obviously important, and neither is trivial to answer. The more "complex" or detailed a model becomes, the

more difficult it is, in general, to understand its behavior; particularly the behavior that is relevant for understanding the behavior of the real system that is being modeled. As the complexity of a model increases, so does our ability to understand what that model itself is really doing.

Example: Tierra

Tierra, developed by Tom Ray of the University of Delaware and the ATR Human Information Processing Research Laboratories in Kyoto, is an extreme bottom-up approach to the simulation of the evolution of artificial organisms at the level of the genome. *Tierra* is designed to provide a "computer-laboratory" environment in which Darwinian evolution can proceed entirely without any intervention from a human operator.

The organisms of *Tierra* are machine-language computer programs consisting of linear strings of an assembly-language-like code written specifically for *Tierra*. A program – or organism – evolves either by mutation or recombination. A typical evolution of a Tierran system starts from a single organism that is capable of self-reproduction. Errors occasionally (and deliberately are made to) creep into the system, rendering some organisms incapable of further self-reproduction and mutating others so that they are able to produce offspring more quickly and efficiently.

The "struggle of evolution" within *Tierra* is essentially a struggle for CPU-time and computer memory. "Survival-of-the-fittest" means that the fittest organisms of the population are those that have managed by whatever means (or by whatever strings of code they have been able to find or construct) to capture more of these available time and space resources than other organisms. Organisms that are able to reproduce quickly and use up relatively little computer memory space in doing so, therefore come to dominate the population.

There are two reasons for mentioning this model:

1. **Simulation versus Instantiation.** *Tierra* illustrates the difference between what Ray calls a *simulation* of artificial life and an *instantiation* of artificial life. In a simulation, computer data structures are explicitly designed to represent real biological entities, whether they be predators and prey, cells, or whatever. In contrast, in an

instantiation of artificial life, computer data structures do not have to explicitly represent any real organism or process. Rather, data structures must obey rules that are abstractly related to the rules governing real processes.

2. Complexity of Instantiation Rivals Complexity of Life.

Tierra appears to capture enough of the dynamics of real Darwinian evolution to harbor some of the same levels of complexity. For example, one typically observes a rich diversity of species in the *Tierran* population, just as one does in nature. There are many interesting examples of basic evolutionary phenomena as well, such as symbiosis and parasitism (though multicellularity has proven more of a challenge to obtain in *Tierra*). The important point to be made here, however, is that because of the inherent, irreducible complexity of *Tierra* "the program," the problem of describing and understanding the behavioral characteristics of *Tierra* "the instantiation of artificial-life" has effectively been rendered to be *just about as complex a problem of describing and understanding the behavioral characteristics of a real biological ecology*. In short, the cost of developing a realistic instantiation of a real complex system is having just as difficult a task of ascertaining what is "really going on" in the instantiation as ascertaining what is really going on in the real system.

Might a similar complexity-explosion await us in modeling land combat? The answer, at this juncture, is unclear. Nonetheless, a great irony possibly awaits us in developing detailed models of land combat. It is entirely conceivable, for example, that the only kind of model that is able to provide us with a deep enough insight into whatever latent high-level patterns exist in combat from which we can conclude anything useful, is a model that is itself fundamentally just as complex a dynamical system as land combat. In such a case, understanding the model becomes just as daunting a task as understanding land combat.

The point worth stressing here, however, is that one must always be mindful of the fact that it is *not just the model of a system that one is after*; one is also after an understanding of what the real system is doing. Being able to develop a sophisticated model of combat – such as by using a system like SWARM – is therefore only part of the story. One must be equally diligent in applying the tool-chest of ways of looking at the behaviors of systems that complex system theory provides. In short, the goal of a complex systems

theory approach to understanding land warfare must not be to develop a model *and stop*, but must include options and strategies for understanding the overall behavior of the model as well.

Thoughtful discussions about the general use of models are given by Denning [17] and Casti [14]. Bankes [3] discusses the advantages and disadvantages of using exploratory modeling for policy analysis.

Assessment of General Applicability

There are several useful ways to organize a list of potential applications of complex systems theory to warfare. The simplest, and most direct, way, is to provide a short discussion of how each and every tool and methodology that goes under the broad rubric of "complexity theory" – including cellular automata, evolutionary programming, fuzzy logic, neural networks, and so on – applies to warfare. The drawback to this approach, of course, is that because of the depth and breadth of available tools, one can soon become hopelessly lost in a meaningless sea of technical jargon. Moreover, the specter of having the overall approach be insipidly, and incorrectly, branded a "solution in search of a problem," is unappealing.

An alternative (and complementary) approach is to start with a list of the most pressing problems associated with land combat – going down the list from predictions of battlefield attrition, to command and control, to fire support, to intelligence, and so on – and to provide suggested avenues of exploration using complex systems theory. A drawback to this approach is that because it starts out with a specific list, it is, in principle, capable only of eliciting the most promising applications to *existing* (i.e. conventional) problems, leaving out what may be the most promising set of applications of complex systems theory to what, in conventional terms, may not (yet!) be recognized as an "issue" or "problem." Indeed, complex systems theory's greatest legacy may prove to be not a set of answers to old questions, but an *entirely new set of questions* to be asked of combat and what really happens on the battlefield.

In light of the above discussion, I decided that the most prudent approach to providing a framework for discussing the possible applications of complex systems theory to land warfare is one that respects both sides of the equation. I define eight separate *Tiers of Applicability*, ranging roughly from least risk and least potential payoff (at least, as far as a practical applicability is concerned) for Tier I, to greatest risk and greatest potential payoff for Tier VIII (see table 3):


- **Tier I:** *General metaphors for complexity in war*
- **Tier II:** *Policy and General Guidelines for Strategy*

- **Tier III:** *Conventional warfare models and approaches*
- **Tier IV:** *Description of the complexity of combat*
- **Tier V:** *Combat technology enhancement*
- **Tier VI:** *Combat aids for the battlefield*
- **Tier VII:** *Synthetic combat environments*
- **Tier VIII:** *Original conceptualizations of combat*

Risk vs. Potential Payoff

Because of the very speculative nature of many of the individual applications making up these eight tiers – applications range from those that are currently undergoing some preliminary form of development to those that currently exist only as vague theoretical possibilities – conventional risk-benefit analysis does not apply. Figure 5, provided for illustrative purposes only, shows very roughly how the applications in each of these eight tiers relates to those in other tiers according to their risk versus potential payoff tradeoff.

Figure 5. Risk versus potential payoff for the eight tiers of applicability (see table 3)



Risk is defined loosely to mean the projected level of research effort that must be invested in order to convert a basic concept into a finished practical application. The greater the degree of investment, in terms of dollars, resources and time, the greater the projected risk. If the projected level of commitment is relatively small – as it is for, say, the first "metaphor" tier of applicability, since it costs effectively nothing to periodically color one's language with "nonlinear metaphors" – the risk is assumed small.

Payoff is defined to mean the potential impact a given application might have on warfare in general and land combat in particular. The greater the potential impact an application might have, the higher is its expected payoff. The *Santa Fe Institute's* SWARM project, for example, when it is completed, has the potential to revolutionize the way complex systems modeling is done in general and therefore shows great promise to significantly impact the way land combat is modeled. Thus, Swarm is unquestionably a "high-payoff" application of complex systems theory.

Table 3. Eight Tiers of Applicability

Tier of Applicability	Description	Examples
<i>I. General Metaphors for Complexity in War</i>	Build and continue to expand base of images to enhance conceptual links between complexity and warfare	nonlinear vice linear synthesist vice analytical edge-of-chaos vice equilibrium process vice structure holistic vice reductionist
<i>II. Policy and General Guidelines for Strategy</i>	Guide formulation of policy and apply basic principles and metaphors of CST to enhance and/or alter organizational structure	Use general metaphors lessons learned from complex systems theory to guide and shape policy making; Use genetic algorithms to evolve new forms
<i>III. "Conventional" Warfare Models and Approaches</i>	Apply tools and methodologies of CST to better understand and/or extend existing models	chaos in Lanchester equations chaos in arms-race models analogy with ecological models
<i>IV. Description of the Complexity of Combat</i>	Describe real-world combat from a CST perspective	power-law scaling Lyapunov exponents entropic parameters
<i>V. Combat Technology Enhancement</i>	Apply tools and methodologies of CST to certain limited aspects of combat, such as intelligent manufacturing, cryptography and data dissemination	intelligent manufacturing data compression cryptography IFF computer viruses fire ants
<i>VI. Combat Aids</i>	Use CST tools to enhance real-world combat operations	autonomous robotic devices tactical picture agents tactics/strategy evolution via GA
<i>VII. Synthetic Combat Environments</i>	Full system models for training and/or to use as research "laboratories"	agent-based models (<i>SimCity</i>) Soar/IFOR SWARM
<i>VIII. Original Conceptualizations of Combat</i>	Use CST-inspired basic research to develop fundamentally new conceptualizations of combat	pattern recognition controlling/exploiting Chaos Universality?

Tier I: General Metaphors for Complexity in War

"If I can't picture it, I can't understand it."

– A. Einstein

"You don't see something until you have the right metaphor to let you perceive it." – Thomas Kuhn

The lowest, but certainly not shallowest, tier of applicability of complex systems theory consists of developing a set of *metaphors* by which war in general, and land combat in particular, can be understood. This set of metaphors represents a new world-view in which the battlefield is seen as a conflict between two self-organizing living-fluid-like organisms consisting of many mutually interacting and co-evolving parts.

What is a Metaphor?

Etymologically, metaphor (the Greek *metafora*, "carry over") means "transfer" or "convey," the transference of a figurative expression from one area to another. According to the 3rd edition of the *American Heritage Dictionary* [2], a metaphor is "a figure of speech in which a word or phrase that ordinarily designates one thing is used to designate another, thus making an implicit comparison. One thing conceived as representing another." The *Encyclopedia Britannica Online*⁷ adds that metaphor "makes a qualitative leap from a reasonable, perhaps prosaic comparison, to an identification or fusion of two objects, to make one new entity partaking of the characteristics of both. Many critics regard the making of metaphors as a system of thought antedating or bypassing logic." In the present context, we can say that the first tier of applicability of complex systems theory to land warfare represents a *reservoir of metaphorical concepts and images* with which land warfare can be illuminated and reinterpreted in a new light.

From the standpoint of the amount of "development time" that is required to make use of metaphor, the risk is effectively zero. One either chooses to color one's language with a particular set of metaphors or one does not. The only groundwork that has to be done is to carefully choose the right set of metaphors. Yet, because of the profound relationship that exists between metaphors and the concomitant reality our language and

⁷

The World-Wide-Web URL address is <http://www.eb.com>.

thoughts create for us, incorporating metaphors borne of complex systems theory into our discussions of combat can potentially radically alter our general understanding of warfare. Therefore, the potential "payoff" is great. But one must at the same time be mindful of using a metaphor, or metaphors, on appropriate levels.

A metaphor can apply either to one particular idea or image that is transferred from, or provides a bridge between, one discipline to another, or, more generally, to a symbolic relation that unites the paradigmatic way of viewing an entire field of knowledge. Emmeche and Hoffmeyer [20] identify four different levels of metaphorical "signification-transfer" in science as follows:

- *Level-1:* The transfer of single terms to other contexts to create new meaning.
- *Level-2:* Construction of analogies as part of a specific theory or a general and systematic inquiry to elucidate phenomena. The analogy may simply be a heuristic device or a component of an apparently final theory.
- *Level-3:* A unifying view of an entire paradigm, often symbolized by a specific term that refers to the whole frame of understanding under a given paradigm.
- *Level-4:* The most comprehensive level of signification is the level on which science itself is understood as irreducibly metaphorical.

While metaphors can be, and often are, misused, they frequently serve as powerful conceptual vehicles by which a set of tools, models and theories is borrowed from one discipline and meaningfully translated to apply to another discipline. For those that say that metaphors are by their nature somehow "shallow" and not "scientific," one need only be reminded that much of science itself advances first by metaphor. Think of Rutherford's analogy of the solar system for the atom or Faraday's use of magnetized iron filings to think about electric fields, among many other examples. The collection of papers edited by Ortony [45] has a section devoted to the significance of metaphor in science.

"Metaphors may be didactic or illustrative devices, models, paradigms, or root images that generate new models. Some metaphors are heuristic, whereas

others constitute new meaning...Borrowing is metaphoric in several ways. Theories and models from other disciplines may sensitize scholars to questions not usually asked in their own fields, or they may help interpret and explain, whether that means a framework for integrating diverse elements or hypothetical answers that cannot be obtained from existing disciplinary resources. When a research area is incomplete, borrowing may facilitate an inductive open-endedness. It may function as a probe, facilitating understanding and enlightenment. Or, it may provide insight into another system of observational categories and meanings, juxtaposing the familiar with the unfamiliar while exposing similarities and differences between the literal use of the borrowing and a new area." [33]

One could argue that much of our reality is structured by metaphor, although we may not always be explicitly aware of this. Lakoff and Johnson [34] suggest that "many of our activities (arguing, solving problems, budgeting time, etc.) are metaphorical in nature. The metaphorical concepts that characterize those activities structure our present reality. New metaphors have the power to create a new reality... we understand a statement as being true in a given situation when our understanding of the statement fits our understanding of the situation closely enough for our purposes." Sometimes the only way to gain a further, or deeper, understanding of an "accepted" reality is to take a step sideways and reinterpret that reality from an alternative vantage point. Creative metaphors help us take that sideways step.

In speaking about the general role that research centers, such as the *Santa Fe Institute*, play in helping decide what metaphors are or are not appropriate for given problems, the economist Brian Arthur argues that [66]

"...the purpose of having a Santa Fe Institute is that it, and places like it, are where the metaphors and a vocabulary are being created in complex systems. So if somebody comes along with a beautiful study on the computer, then you can say 'Here's a new metaphor. Let's call this one *the edge of chaos*,' or whatever. So what the Santa Fe Institute will do, if it studies enough complex systems, is to show us the kinds of patterns we might observe, and the kinds of metaphor that might be appropriate for systems that are moving and in process and complicated, rather than the metaphor of clockwork."

Metaphors and War

Though conventional military thinking has, through history, been arguably dominated by the clockwork precision of the "Newtonian" metaphor – exemplified by the often cited view of combat between two adversaries as "a collision between two

billiard-balls" – one can find examples of complexity-ridden metaphors in many important military historical writings.

Consider Sun-Tzu's analogy of a military force to water [64]:

"So a military force has no constant formation, water has no constant shape: the ability to gain victory by changing and adapting according to the opponent is called genius."

Here Sun-Tzu likens movement and maneuver on the battlefield to the complex dynamics of fluid flow, which is a very apt metaphor for "combat as a complex system." He also underscores the importance of adaptability on the battlefield, which is the hallmark of any healthily evolving complex adaptive system.

In a recent article in *International Security*, entitled "Nonlinearity and Clausewitz," Beyerchen argues persuasively that much of Clausewitz's military thought was colored by a deep intuitive understanding of nonlinear dynamics [8]:

"*On War* is suffused with the understanding that every war is inherently a nonlinear phenomenon, the conduct of which changes its character in ways that cannot be analytically predicted. I am not arguing that reference to a few of today's 'nonlinear science' concepts would help us clarify confusion in Clausewitz's thinking. My suggestion is more radical: in a profoundly unconfused way, he understands that seeking exact analytical solutions does not fit the nonlinear reality of the problems posed by war, and hence that our ability to predict the course and outcome of any given conflict is severely limited."

Clausewitz's "fog-of-war," "center-of-gravity" and "friction," of course, are well known. In the last section of Chapter 1, Book One, Clausewitz compares war to a "remarkable trinity" composed of (1) the natural force of hatred among the masses, (2) war's inherent element of chance, and (3) war's subordination to governmental policy. He concludes with a wonderful visual metaphor that anticipates one of the prototypical experimental demonstrations of deterministic chaos: "Our task therefore is to develop a theory that maintains a balance between these three tendencies, like an object suspended between three magnets." [8]

In another section, Clausewitz takes a bold stride beyond the "combat as colliding billiard-balls" metaphor, and anticipates almost directly the core element of the new "combat as complex adaptive systems" view:

"...war is not an exercise of the will directed at inanimate matter, as in the case with the mechanical arts, or at matter which is animate but passive and yielding as in the case with the human mind and emotion in the fine arts. In war, the will is directed at an *animate object that reacts*."

It is a great testament to Clausewitz's brilliance and deep insight that he was able to recognize and exploit such provocative imagery to illustrate his ideas, insofar as there was no such field as "nonlinear dynamics" in his day.

In contrast, metaphors of nonlinearity are today much more commonplace, thanks in large part to the popularization of such "new sciences" as nonlinear dynamics, deterministic chaos and complex systems theory. To the extent that Clausewitzian theory itself accurately describes the fundamentals of war, the metaphors borne of nonlinear dynamics and complex systems theory therefore also have much to tell us. However, one should at the same time be cautious of "plumbing the wells of metaphor" too deeply, or of expecting, *free-of-charge*, a greater clarity or eloquence of expression in return. An unbridled, impassioned use of metaphor alone, without taking the time to work out the details of whatever deeper insights the metaphor might be pointing to, runs the risk of both shallowness and loss of objectivity.

Having said this, it is still true that if, in the end, it turns out that complex systems theory provides no genuinely new insights into war *other than to furnish a rich scaffolding of provocative and suggestive metaphors* around which an entirely new view of warfare can be woven – i.e. if the "signification-depth" is essentially confined to levels 1 and 2 of Emmeche's and Hoffmeyer's hierarchy (see above) – complex systems theory will have nonetheless fulfilled an enormously important function. Time will certainly tell if the new metaphors are as deep and meaningful as they at first appear or are "just another passing fad" that will soon fade from view. But just the fact alone that these new metaphors are being actively engaged in serious discussion at high levels⁸ is enough to suggest that the consensus reality is already being altered. In a very real sense, the reality of "war as complex adaptive system"

⁸ As an example, a recent conference sponsored by the Marines Corps Combat and Development Command – entitled *Non-Linear Studies and Their Implications for the US Marine Corps* – and at which the interim results of this project were briefed, attracted more than 100 participants. As another example, Tom Czerwinski's intensive but popular day-long seminars on the military implications of the "new sciences," conducted at the *National Defense University* in Washington, D.C., regularly draw many high-ranking leaders from all branches of the military.

did not exist before the discussion started. And the longer a serious discussion continues in earnest, the longer the participants will have to develop a more meaningful complexity-metaphor ridden vocabulary of combat, and the deeper and more compellingly the images and concepts can take root in their minds. Indeed, if these new metaphors capture anything at all that is basic to war, they will, in time, inevitably take just as firm a hold of the language of war for future generations as the Newtonian metaphor of "colliding billiard-balls" has taken hold of the military thinking of past generations.

Metaphor Shift

The first tier of applicability of complex systems theory to warfare consists of a set of new metaphors by which war in general, and land combat in particular, can be understood. This set of metaphors represents a shift, *away* from the old "Newtonian" word-view – that emphasizes equilibrium and sees the battlefield as an arena of colliding objects obeying simple, linear laws and possessing little or no internal structure – *to* a new (but, ironically, older) "Heraclitian"⁹ world-view, that emphasizes process and sees the battlefield as a conflict between two self-organizing living-fluid-like organisms consisting of many mutually interacting and co-evolving parts.

The new "Heraclitian" metaphor is a rich interlacing tapestry of ideas and images, woven of five basic conceptual threads – *nonlinearity*, *deterministic chaos*, *complexity*, *self-organization* and *emergence*:

- **Nonlinearity.** In colloquial terms, nonlinearity refers to the property that the whole is not necessarily equal to the sum of its parts. More precisely, if f is a nonlinear function or operator, and x is a system input (either a function or variable), then the effect of adding two inputs, x_1 and x_2 , first and then operating on their sum is, in general, not equivalent to operating on two inputs separately and then adding the outputs together; i.e. $f(x+y)$ is, in general, not equal to $f(x) + f(y)$. For example, for the nonlinear

⁹ The label "Heraclitian," as here used, is patterned after the label used by the geneticist Richard Lewontin to refer to scientists who see the world as a process of flow [35]. Heraclitus was an Ionian philosopher who argued that the world is in a constant state of flux. One of his more famous passages is, "You can never step into the same river twice."

function $f(x) = x^2$, $f(x_1 + x_2)^2 = (x_1 + x_2)^2 = x_1^2 + x_2^2 + 2x_1x_2$, and the last term $- 2x_1x_2 -$ appears as an additional quantity. Nonlinear systems can therefore display a disproportionately small or large output for a given input. Nonlinear systems are also generally very difficult to deal with mathematically, which is the main reason why they are usually replaced by linear approximations. While the act of linearization simplifies the problem, however, most of the interesting behavior of the real nonlinear system is washed away in the process. The "Heraclitian" metaphor reminds us that *war is inherently nonlinear*, and, as such, ought not be "linearized" away in an attempt to achieve a simplified "solution."

- **Deterministic Chaos.** Deterministic chaos refers to irregular or chaotic motion that is generated by nonlinear systems evolving according to dynamical laws that uniquely determine the state of the system at all times from a knowledge of the system's previous history. It is important to point out that the chaotic behavior is due neither to external sources of noise nor to an infinite number of degrees-of-freedom nor to quantum-mechanical-like uncertainty. Instead, the source of irregularity is the exponential divergence of initially close trajectories in a bounded region of the system's phase space. This sensitivity to initial conditions is sometimes popularly referred to as the "butterfly effect," alluding to the idea that chaotic weather patterns can be altered by a butterfly flapping its wings. A practical implication of chaos is that its presence makes it essentially impossible to make any long-term predictions about the behavior of a dynamical system: while one can in practice only fix the initial conditions of a system to a finite accuracy, their errors increase exponentially fast. This does not mean, however, that short-term predictability is lost, for deterministic chaos also implies that within what appears to be erratic motion lies an underlying order. This underlying order can potentially be exploited to make short-term predictions.
- **Complexity.** Complexity is an extremely difficult "I know it when I see it" concept to define, partly because its real meaning consists of two parts, neither of which is easy to quantify: (1) *system complexity*, which refers to the structural, or organizational, complexity of a system – examples include the interacting molecules in a fluid and the

network topology of neurons in a brain; and (2) *behavioral complexity*, which refers to the complexity of actual behavioral patterns exhibited by complex, or simple, systems, as they evolve – examples include deterministic chaos, multistability, multifractality, and so on. A succinct summary of what nonlinear dynamics and complex systems theory together show, conceptually, is that it is possible to *both* (a) reduce the system complexity to only a relatively few degrees-of-freedom (i.e. to generate simplicity from complexity), *and* (b) have simple low-dimensional dynamics exhibit complex behaviors (i.e. to generate complexity from simplicity). Both of these strategies must be used in dealing with, and understanding, the complexities of combat.

- **Self-Organization.** Self-organization is the spontaneous emergence of macroscopic nonequilibrium organized structure due to the collective interactions among a large assemblage of simple microscopic objects. Patterns emerge spontaneously when, say, certain environmental factors change. It is important to understand that these self-organized patterns arise out of a purely internal dynamics, and not because of any external force. It used to be believed that any kind of order must be due to some "oracle" imposing the order from outside of the system. Self-organization shows that no such oracle is needed. Examples of self-organization abound: convection flows in fluids, morphogenesis in biology, concentration patterns in chemical reactions, atmospheric vortices, *etc.*
- **Emergence.** Emergence refers to the appearance of higher-level properties and behaviors of a system that – while obviously originating from the collective dynamics of that system's components – are neither to be found in nor are directly deducible from the lower-level properties of that system. Emergent properties are properties of the "whole" that are not possessed by any of the individual parts making up that whole. Individual lines of computer code, for example, cannot calculate a spreadsheet; an air molecule is not a tornado; and a neuron is not conscious. Emergent behaviors are typically novel and unanticipated.

The metaphor shift, of course, involves many more concepts and images than that to which these five "conceptual threads" alone attest, though these five certainly represent the core set. For

example, in general terms, one can say that where the old metaphor stressed *analysis*, in which a system is understood by being systematically broken down into its parts, the new metaphor stresses *synthesis*, in which how a system behaves is discovered by building it up from its pieces. Where the old metaphor stressed a *mechanistic dynamics*, in which combat is viewed as a sequence of strictly materially caused events, the new metaphor stresses an *evolutionary dynamics*, in which combat is viewed as a co-evolutionary process among adapting entities. Where the old metaphor stressed *equilibrium* and *stability*, in which "solutions" to combat are found after the system "settles down," the new metaphor stresses the importance of *far-from-equilibrium* states and the continual quest for *perpetual novelty*, in which combat never settles down to an equilibrium, and is always in pursuit of the so-called "edge-of-chaos" (see Part I, page [76]).

Tables 4 and 5 illustrate the metaphor shift from the old "Newtonian" view of the combat to the new "Heraclitian" view. Table 4 compares the old and new metaphors from the standpoint of their respective vocabularies. It is by no means a complete list of relevant words and concepts, and is meant only to capture some of the essential ideas. Table 5 extends this list of words by using various contexts within which to compare some of the basic principles underlying the two metaphors.

Further speculations on possible connections between the "new sciences" and war on a metaphor level can found in Beyerchen [8], Beaumont [6], Hedgepeth [25], Saperstein [54], and Schmitt [54].

Table 4. The shift from "Newtonian" to "Heraclitian" metaphors as reflected in their corresponding *Vocabularies* of concepts and images

Old Metaphor ("Newtonian")	New Metaphor ("Heraclitian")
Analytical	Synthesist
Basic elements are "Quantities"	Basic elements are "Patterns"
Behavior is Contingent and Knowable	Behavior is Emergent and often Unexpected
Being	Becoming
Clockwork Precision	Open-ended Unfolding
Closed System	Open System
Complexity Breeds Complexity	Complexity Can Breed Simplicity
Deterministic	Deterministically Chaotic
Equilibrium	Far-From-Equilibrium/ "perpetual novelty"
Individualistic	Collective
Linear	Nonlinear
Linear Causation	Feedback Loop/Circular Causality
Mechanistic Dynamics	Evolutionary Dynamics
Military "Operation"	Military "Evolution"
Combat as Collision Between Newtonian "Billiard-Balls"	Combat as Self-Organized ecology of living "fluids"
Order	Inherent Disorder
Predesigned	Emergent
Predictable	Unpredictable
Quantitative	Qualitative
Reductionist	Holistic
Solution	Process and Adaptation
Stability	"Edge-of-Chaos"
Top-Down	Bottom-Up <i>and</i> Top-Down
<i>etc.</i>	<i>etc.</i>

Table 5. A Comparison Between some of the principles underlying "Newtonian" and "Heraclitian" Metaphors

Context	"Newtonian" Metaphor	"Heraclitian" Metaphor
<i>Complex behavior</i>	Complex behavior requires complex models	Simple models often suffice to describe complex systems; <i>complexity from simplicity and simplicity from complexity</i>
<i>Patterns of behavior</i>	Each qualitatively different pattern of behavior requires a different equation	Qualitatively different patterns of behavior can be described by the same underlying equation
<i>Description of Behavior</i>	Each qualitatively different kind of behavior requires new equation or set of equations	One equation harbors a multitude of qualitatively different patterns of behavior
<i>Effects of small perturbations</i>	Small perturbations induce small changes	Small perturbations can have large consequences
<i>How to understand system</i>	A system can be understood by breaking it down into and analyzing its simpler components	Systems can be understood only by respecting the mutual interactions among its components; look at the <i>whole system</i>
<i>Origin of Disorder</i>	Disorder stems mainly from unpredictable forces <i>outside</i> of system	Disorder can arise from forces entirely <i>within</i> the system
<i>Origin of Order</i>	Order must be imposed from <i>outside</i> the system	Order can arise in a purely self-organized fashion <i>within</i> the system
<i>Nature of observed order</i>	Order, once present, is pervasive and appears both locally and globally	A system may appear locally disordered but possess global order
<i>"Goal"</i>	Goal is to develop "equations" to describe behavior; determined by isolating effect of one variable at a time	Goal is to understand how entire system responds to various contexts, with no one variable dominating
<i>Type of "solutions"</i>	Goal is to search for "optimal" solution	No optimal solution exists, as the set of problems and constraints continuously changes
<i>Predictability</i>	Assuming that the "correct" model is found and initial conditions are known exactly, everything is predictable and controllable,	Long-term predictability may be unattainable even in principle; behavior may be predicted for short-times only
<i>Nature of causal flow</i>	Causation flows from the <i>bottom up</i>	Causation flows both from <i>bottom up and from the top down</i>

Tier II: Policy and General Guidelines for Strategy

What Does the New Metaphor Give Us?

Roger Lewin, in his popular book on complexity [37], reproduces a fragment of a conversation he once had with Patricia Churchland, a well-known neurobiologist:

"Is it reasonable to think of the human brain as a complex dynamical system?" I asked. "Its obviously true," she replied quickly. "But so what? Then what is your research program? ... What research do you do?"

Notice that if "human brain" is replaced in this fragment by "land combat," the fragment retains the potent sting of Churchland's challenge. Every new research endeavor must begin with at least these two basic steps: (1) a prior justification that the endeavor is a reasonable one to consider undertaking, and (2) a plan of attack. As far as the endeavor of applying complex systems theory to land warfare is concerned, the first step is easy: land combat – *on paper* – has almost all of the key attributes that any reasonable list of attributes of a complex adaptive system must include (see table 1 appearing in the **Executive Summary**). The second step is by far the more difficult one to take: now that we have established the similarity, *what do we do with the connection?*

Policy

"If you have a truly complex system, then the exact patterns are not repeatable. And yet there are themes that are recognizable. In history, for example, you can talk about 'revolutions,' even though one revolution might be quite different from another. So we assign metaphors. It turns out that an awful lot of policy-making has to do with finding the appropriate metaphor. Conversely, bad policy-making almost always involves finding inappropriate metaphors. For example, it may not be appropriate to think about a drug 'war,' with guns and assaults."¹⁰

The first step to take beyond merely weaving threads of metaphor is to apply the basic lessons learned from the study of complex systems to how we formulate strategy and general policy. This assumes implicitly that political systems and world communities can be just as well described as complex adaptive systems as can be human brains and collections of combat forces. For example, consider that the essence of a (successfully

¹⁰ Quote attributed to Brian Arthur in reference [66].

evolving) complex adaptive complex is to exist in a far-from-equilibrium state and to continually search for novelty and new solutions to changing problems. An important lesson learned for a complex systems theoretic approach to policy making is therefore to shift from general policies that emphasize a *means to achieve stability* to policies that encourage a *continual co-evolution of all sides*.

Similarly, nonlinear dynamics teaches us that it is the nonlinearities embedded in a set of processes that are responsible for instability and irregular appearing behavior. The lesson learned here for dealing with adversaries in a conflict is to focus attention on the nonlinear drivers of an enemy's system, for these are the elements that can potentially induce the greatest effect from the least effort.

Table 6 provides general behavioral guidelines and strategies for conduct and policy making derived from the "Heraclitian" metaphor.

Organizational Structure

A thorough understanding of complexity and complex adaptive systems can be applied to enhance and/or alter organizational and command structures. Practical techniques can be developed for the military to re-examine their metaphors and beliefs, and to adapt new ones as conditions change. General techniques for building "learning organizations," such as adapted from a systems theory model of management and described in Senge's Fifth Discipline [57] can be applied. These techniques include general strategies for dealing with the unexpected and/or accidental, and resolving the dichotomous needs for both stability and creativity. The overall objective is to use the basic lessons of complex systems theory to develop sets of internal organizational "rules" and strategies that are more conducive to adaptation and self-organization. Genetic algorithms, too, can be used to search for better command and control structures (see *Tier-VI*).

Table 6. General behavioral guidelines and strategies for conduct and policy making derived from the "Heraclitian" metaphor

General Guideline	Description
<i>Exploit Collective Action</i>	Exploit the synchronous parallel cooperative effort of many low-level agents rather than solving for a global solution all at once
<i>Expect Change</i>	Never take your eyes off, or turn your back on, a system – systems continually evolve and change
<i>Stop looking for "Optimal" Solutions</i>	Forget about optimizing a solution to a problem – the problem is constantly changing
<i>"Guide" Behavior, Do not "Fix It"</i>	Emphasize the process vice solution approach. Instead of focusing on "single points" of a trajectory – or snapshots in time of key events that unfold as a policy is implemented – focus on how to continue to nudge the system in a favorable direction
<i>Look for Global Patterns</i>	Search for global patterns in time and/or space scales higher than those on which the dynamics is defined; systems can appear locally disordered but harbor a global order
<i>Apply Holistic Understanding</i>	Focus more on identifying interdependent behaviors (i.e. how a system responds to different contexts and when interdependent sets of parameters are allowed to change) rather than looking for how a system changes when everything is left constant and one parameter at a time is allowed to change
<i>Focus More "Within" for Understanding Source of Apparent Irregularity</i>	Irregular and random appearing behavior that appears to be due either to outside forces or elements of chance may be due solely to the internal dynamics of a system
<i>Look for Recognizable Themes and Patterns</i>	Exact patterns may not be repeated but the general underlying themes will remain the same
<i>Focus on Process vice Static Measures</i>	Study the logic, dynamics, process, etc. not the material constituents of a system
<i>To "Break Down" Does Not Always Mean to "Get Simpler"</i>	If the behavior of a system is described by a fractal, successively finer views of the fractal reveal successively finer levels of detail. Things do not necessarily "get simpler."
<i>Exploit Nonlinearity</i>	Focus attack on the nonlinear processes in an enemy's system, for these are the processes that can potentially induce the greatest effect from least effort
<i>Do not Necessarily Frown on Chaos</i>	A bit of irregularity or "chaos" is not necessarily a bad thing, for it is when a system is at the "edge-of-chaos" that it is potentially best able to adapt and evolve
<i>Exploit Decentralized Control</i>	Encourage decentralized control, even if each "patch" attempts to optimize for its own selfish benefit; but maintain interaction among all patches
<i>Find Ways to be More Adaptable</i>	The most "successful" complex systems do not just continually adapt, they struggle to find ways to adapt better; move towards a direction that gives you more options

Intelligence Analysis

Conventional intelligence analysis consists of first assessing the information describing a situation and then predicting its future development. The task is complicated by the fact that the available information is often incomplete, imprecise and/or contradictory. Moreover, the information may be falsified or planted by the adversary as part of a deliberate disinformation campaign. The traditional reductionist approach of dealing with these problems consists of six general steps [67]:

1. *Data Management*: all collected information is first processed to conform to selected forms of data management (index cards, computer data-bases, etc.)
2. *Reliability Grading*: information is graded for its veracity, which depends on such factors as source of information and existence of collaborative sources
3. *Subject Sorting*: information is broken down into more manageable parts, typically by subject
4. *Relevance Filtering*: subject sorted information is parsed for degrees of relevance (sorting "wheat from chaff")
5. *Search for "Trigger Facts"*: filtered information is searched for characteristic facts and/or events that are known or suspected as being triggers or indicators of specific future events
6. *Search for Patterns*: information is examined for clues of patterns of activity and linkages to assist in making specific predictions

While there is nothing sacrosanct about any of these six steps, and each intelligence analyst undoubtedly evolves his or her own unique style and approach, the fundamentally reductionist manner in which all such analysis is invariably conducted suffers from a number of significant drawbacks [67]. For example, the process of collating the information often curtails an analyst's ability to respond quickly to important indicators. Adherence to a predefined order (such as requiring that all information be fit into an existing data management system) may also lead to difficulties in assimilating any unexpected or unusual

information. In the worst case, information that does not strictly conform to accepted or understood patterns and categories or does not conform to an anticipated course of events, may be ignored and/or discarded. Finally, a predisposition to filtering out and assimilating only "conventional" (i.e. "doctrinal") forms of information makes it hard for an analyst to appreciate (and therefore respond to) other, perhaps unconventional, variables that may in fact play a vital role in determining the future behavior of a system (and without which an accurate prediction may be impossible to obtain).

Complex systems theory, with its emphasis on pattern recognition and its general openness-of-mind when it comes to what variables and/or parameters might be relevant for determining the future evolution of a system, has many potentially useful suggestions to offer the intelligence analyst for his/her analysis of raw intelligence data. For example, complex systems theory persuades an analyst, in general, not to discard information solely on the basis of that information not conforming to a "conventional wisdom" model of an adversary's pattern of activity. Instead, and as has been repeatedly stressed throughout this paper, complex systems theory teaches us to recognize the fact that apparently irrelevant pieces of information may contain vital clues as to an adversary's real intentions.

Policy Exploitation of Characteristic Time Scales of Combat

A fundamental property of nonlinear systems is that they generally react most sensitively to a special class of aperiodic forces. Typically, the characteristic time scales of the optimal driving force match at all times the characteristic time scales of the system. In some cases the optimal driving force as well as the resulting dynamics are similar to the transients of the unperturbed system [27].

The information processing in complex adaptive systems and the general sensitivity of all nonlinear dynamical systems to certain classes of aperiodic driving forces are both potentially exploitable features. Recall that one of the distinguishing characteristics of complex systems is their information processing capability. Agents in complex adaptive systems continually sense and collect information about their environment. They then base their response to this information by using internal models of the system, possibly encoding and storing data about novel situations

for use at a later time. According to the *edge-of-chaos* idea (see Part I [28], page 76), the closer a system is to the edge-of-chaos – neither too ordered nor too chaotic – the better it is able to adapt to changing conditions. In Kauffman's words,

"Living systems exist in the...regime near the edge of chaos, and natural selection achieves and sustains such a poised state...Such poised systems are also highly evolvable. They can adapt by accumulation of successive useful variations precisely because damage does not propagate widely...It is also plausible that systems poised at the boundary of chaos have the proper structure to interact with and internally represent other entities of their environment. In a phrase, organisms have internal models of their worlds which compress information and allow action...Such action requires that the world be sufficiently stable that the organism is able to adapt to it. Were worlds chaotic on the time scale of practical action, organisms would be hard pressed to cope."¹¹

Now compare this state-of-affairs with Retired USAF Colonel John Boyd's *Observe-Orient-Decide-Act* (OODA) loop [10]. In Boyd's model, a system responds to an event (or information) by first observing it, then considering possible ways in which to act on it, deciding on a particular course of action and then acting. From a military standpoint, both friendly and enemy forces continuously cycle through this OODA process. The objective on either side is to do this more rapidly than the enemy; the idea being that if you can beat the enemy to the "punch" you can disrupt the enemy's ability to maintain coherence in a changing environment. One can also imagine exploiting the relative phase relationship between friendly and enemy positions within the OODA loop. For example, by carefully timing certain actions, one can effectively slow an enemy's battle-tempo by locking the enemy into a perpetual *Orient-Orient* mode.

Cooper [15] has generalized this notion to what he calls "phase-dominance," where the idea is to exploit the natural operating cycles and rhythms of enemy forces and execute appropriate actions exactly when they are needed. In phase-dominance, "time becomes the critical determinant of combat advantage."

Tier III: "Conventional" Warfare Models and Approaches

Tier III consists of applying the tools and methods of nonlinear dynamics and complex systems theory to more or less "conventional models" of combat. The idea on this tier is not so much to develop entirely new formulations of combat so much as extending and generalizing existing forms using a new mathematical arsenal of tools. Examples include looking for chaos in various generalized forms of Lanchester equations, applying nonlinear dynamics to arms-race models, exploiting common themes between equations describing predator-prey relations in natural ecologies and the equations describing combat, and so on.

Testing for the Veracity of Conventional Models

A very practical application of one of the most widely used tools of complex systems theory – namely, *genetic algorithms* – is to the sensitivity analysis and general testing of conventional models of complex systems. Consider, for example, Miller's *Active Nonlinear Test* (ANT) approach to testing the veracity of complex simulation models [40].

As large-scale computational models grow in popularity because of their ability to help analyze critical scientific, military, and policy issues, the same conditions that make them so appealing are also the ones that make testing such models more and more difficult. Such models typically deal with enormously large search, or "solution," spaces and are characterized by a high degree of nonlinearity. Traditional "sensitivity analysis" techniques, which probe for a model's reaction to small perturbations in order to get a feel for how sensitive the model is to variations in values of key control parameters, require simple *linear* relationships within the model in order to be effective.

This last point is a very important one. If the underlying relationships among a model's key parameters is *inherently nonlinear* – as it must be in any reasonably realistic model of a real complex system – then information about the effect of systematically perturbing individual parameters may not be useful in determining the effects of perturbing groups of parameters. To see this, suppose that the "model" is given by the functional form $f(x_1, x_2) = x_1 x_2$ (which is obviously nonlinear). Then $f(x_1 + \Delta x_1, x_2) = x_1 x_2 + x_2 \Delta x_1 = f(x_1, x_2) + x_2 \Delta x_1$. Similarly,

$f(x_1, x_2 + \Delta x_2) = f(x_1, x_2) + x_1 \Delta x_2$. But $f(x_1 + \Delta x_1, x_2 + \Delta x_2) = f(x_1, x_2) + x_2 \Delta x_1 + x_1 \Delta x_2 + \Delta x_2 \Delta x_1$. Thus the effect of changing both x_1 and x_2 simultaneously differs from the adding the effects of the individual perturbations to x_1 and x_2 by the last term, $\Delta x_2 \Delta x_1$. (The linear approximation works well enough, of course, as long as either the perturbations are kept small or the nonlinearity is small.)

The idea behind Miller's ANT approach is to use a genetic algorithm (or any other nonlinear optimization algorithm) to search the space of *sets* of reasonable model perturbations. The objective is to maximize the deviation between the original model's prediction and that obtained from the model under the perturbations. Note that while one could, in principle, detect nonlinearities by exhaustively searching through the space of all possible combinations of pertinent parameters, the potentially enormous space that the resulting combinatoric explosion gives rise to makes such an exhaustive search unfeasible even when only a relatively few parameters are involved. Thus, Miller's objective is to use a genetic algorithm to perform a directed search of groups of parameters.

ANTs work essentially by probing for weakness in a model's behavior. The idea is to obtain an estimate of the maximum error that is possible in a model by actively seeking out a model's worst-case scenarios. Miller is quick to point out that this approach has two limitations: (1) it fails to give an estimate of the likelihood that the worst-case scenarios will actually occur (though other techniques, such as Monte Carlo methods, can be used for this), and (2) the inability to "break" a model by probing its worst-case scenarios does not guarantee a model's overall quality (since a not terribly well designed model could simply be insensitive to its parameters).

Variations of the basic ANT technique could prove useful for testing many existing models and simulations of land combat.

Non-Monotonicities and Chaos

A fundamental lesson of nonlinear dynamics theory is that one can almost always expect to find some manifestation of chaos whenever nonlinearities are present in the underlying dynamics of a model. This fundamental lesson has potentially significant implications for even the simplest combat models.

Miller and Sulcoski [41]-[42], for example, report fractal-like properties and a sensitivity-to-initial conditions in the behavior of a discretized model of the Lanchester equations (augmented by nonlinear auxiliary conditions such as reinforcement and withdrawal/ surrender thresholds).

A recent RAND study¹² has uncovered chaotic behavior in a certain class of very simple combat models in which reinforcement decisions are based on the state of the battle. The study looked at non-monotonicity and chaos in combat models, where "monotonic behavior" is taken to mean a behavior in which adding more capabilities to only one side leads to at least as favorable an outcome for that side.

The presence of nonmonotonicities has usually been interpreted to mean that there is something wrong in the model that needs to be "fixed" and has been either treated as an anomaly or simply ignored. The main thrust of the RAND report is that, while non-monotonicities often do arise from questionable programming skills, there is a source of considerably more problematic non-monotonicities that has its origins in deterministic chaos .

The RAND study found that "a combat model with a single decision based on the state of the battle, no matter how precisely computed, can produce non-monotonic behavior in the outcomes of the model and chaotic behavior in its underlying dynamics."

The authors of the report draw four basic lessons from their study:

- models may not be predictive, but are useful for understanding changes of outcomes based on incremental adjustments to control parameters
- scripting the addition of battlefield reinforcement (i.e. basing their input on time only, and not on the state of the battle) generally eliminates chaotic behavior
- one can identify input parameters figuring most importantly in behavior of non-monotonicities – these are the size of reinforcement blocks and the total number of reinforcements available to each side

¹² J. A. Dewar, J. J. Gillogly and M. L. Juncosa, "Non-Monotonicity, Chaos, and Combat Models," RAND , R-3995-RC, 1991.

- Lyapunov exponents are useful to evaluate a model's sensitivity to perturbations

In general, the RAND report concludes that

"In any combat model that depends for its usefulness on monotonic behavior in its outcomes, modeling combat decision based on the state of the battle must be done very carefully. Such modeled decisions can lead to monotonic behavior and chaotic behavior and the only sure ways (to date) to deal with that behavior are either to remove state dependence of the modeled decisions or to validate that the model is monotonic in the region of interest."

Minimalist Modeling

Dockery and Woodcock, in their massive treatise *The Military Landscape* [18], provide a detailed discussion of many different "minimalist models" from the point of view of catastrophe theory and nonlinear dynamics. Minimalist modeling refers to "the simplest possible description using the most powerful mathematics available and then" adds layers "of complexity as required, permitting structure to emerge from the dynamics." Among many other findings, Dockery and Woodcock report that chaos appears in the solutions to the Lanchester equations when modified by reinforcement. They also discuss how many of the tools of nonlinear dynamics can be used to describe combat.

Using generalized predator-prey population models to model interactions between military and insurgent forces, Dockery and Woodcock illustrate (1) the set of conditions that lead to a periodic oscillation of insurgent force sizes, (2) the effects of a limited pool of individuals available for recruitment, (3) various conditions leading to steady state, stable periodic oscillations and chaotic force-size fluctuations, and (4) the sensitivity to small changes in rates of recruitment, disaffection and combat attrition of simulated force strengths.

This kind of analysis can sometimes lead to counter-intuitive implications for the tactical control of insurgents. In one instance, for example, Dockery and Woodcock point out that cyclic oscillations in the relative strengths of national and insurgent forces result in recurring periods of time during which the government forces are weak and the insurgents are at their peak strength. If the government decides to add too many resources to strengthen its forces, the chaotic model suggests that

the cyclic behavior will tend to become unstable (because of the possibility that disaffected combatants will join the insurgent camp) and thus weaken the government position. The model instead suggests that the best strategy for the government to follow is to use a moderately low level of military force to contain the insurgents at their peak strength, and attempt to destroy the insurgents only when the insurgents are at their weakest force strength level of the cycle.¹³

Generalizations of Lanchester's equations

In 1914, Lanchester introduced a set of coupled ordinary differential equations as models of attrition in modern warfare. The basic idea behind these equations is that the loss rate of forces on one side of a battle is proportional to the number of forces on the other. In one form of the equations, known as the *directed-fire* (or *square-law*) model, the Lanchester equations are given by the linear equations $dR(t)/dt = -\alpha_B B(t)$ and $dB(t)/dt = -\alpha_R R(t)$, where $R(t)$ and $B(t)$ represent the numerical strengths of the red and blue forces at time t , and α_R and α_B represent the constant effective firing rates at which one unit of strength on one side causes attrition of the other side's forces. An encyclopedic discussion of the many different forms of the Lanchester equations is given by Taylor ([62], [63]).

While the Lanchester equations are particularly relevant for the kind of static trench warfare and artillery duels that characterized most of World War I, they are too simple and lack the spatial degrees-of-freedom needed to realistically model modern combat. The fundamental problem is that they idealize combat much in the same way as Newton's laws idealize the real chaos and complexity ridden physics of the world. Likewise, almost all Lanchester equation based attrition models of combat suffer from many basic shortcomings:

- determinism, whereby the outcome of a battle is determined solely as a function of the initial conditions, without regard for Clausewitz's "fog of war" and "friction"
- use of effectiveness coefficients that are constant over time
- static forces

¹³

Reference [77], pages 137-138.

- homogeneous forces with no spatial variation
- no combat termination conditions
- the assumption that target acquisition is independent of force levels
- no consideration of the suppression effects of weapons
- *and so on ...*

Perhaps the most important shortcoming of virtually all Lanchester equation based models is that such models rarely, if ever, take into account the human factor; i.e. the psychological and/or decision-making capability of the individual combatant.

Adaptive Dynamic Model of Combat

The adaptive dynamic model of combat is a simple analytical generalization of Lanchester's equations of combat that adds a basic behavioral dimension by building in a feedback between troop movement and attrition. It is discussed by Epstein [22].

Epstein introduces two new parameters: (1) α_a , which is the daily attrition rate the attacker is willing to suffer in order to take territory, and (2) α_d , which is the daily attrition rate the defender is willing to suffer in order to hold territory. He uses these parameters to express some simple expectations of human behavior. If the defender's attrition rate is less than or equal to α_d , for example, the defender is assumed to remain in place; otherwise this "pain threshold" is exceeded and he withdraws to restore his attrition rate to more acceptable levels. Similarly, if the attackers "pain threshold" is exceeded, he cuts off the attack. Combat is seen as the interplay of "two adaptive systems, each searching for its equilibrium, that produces the observed dynamics, the actual movement that occurs and the actual attrition suffered by each side." [22] Postulating some simple functional forms to express intuitive relationships that must hold true among prosecution, withdrawal and attrition rates, Epstein derives expressions for adaptive prosecution and withdrawal rates for attacking and defending forces. Though we will not go into the details here, Epstein's simple model seems to capture some of the basic behavioral characteristics that are so glaringly missing from Lanchester's equations.

Lotka-Volterra Equations

Studies in predator-prey interactions in natural ecologies have a rich analytical history dating back to the middle 1920s. Around that time, Lotka and Volterra independently proposed the first mathematical model for the predation of one species by another to explain the oscillatory level of certain fish in the Atlantic. If $N(t)$ is the prey population and $P(t)$ is the predator population at time t then $dN/dt = N(a - bP)$, $dP/dt = P(cN - d)$, where a, b, c , and d are positive constants. The model assumes: (1) prey in absence of predation grows linearly with N ; (2) predation reduces prey's growth rate by a term proportional to the prey and predation populations; (3) the predator's death rate, in the absence of prey, decays exponentially; (4) the prey's contribution to the predator's growth rate is proportional to the available prey as well as to the size of the predator population.

What is interesting about these simple Lotka-Volterra equations is that they describe *exactly* the same model as the one Lanchester used to represent land combat. The same kind of oscillatory behavior found in Lanchester's equations, for example, is exhibited by predator-prey systems. Much work, of course, has been done since Lotka's and Volterra's time to generalize their basic equations, including the addition of nonlinear terms to model real-world interactions better, incorporating the complexities of real-world life-cycles and the immune response of hosts in host-parasite systems, modeling interactions between predator-prey systems and their natural environments, exploring the origins of multistability, and so on. However, despite the many conceptual advances that have been made, which today also include the use of sophisticated computer modeling techniques such as multi-agent based simulations, this rich history of analytical insights into the behavior of predator-prey systems has heretofore been largely ignored by conventional operations research "analysis" of combat. Simple Lotka-Volterra-like models of ecologies make up a sizable fraction of the models used in complex systems theory and can potentially be exploited to provide insights into the general behavioral patterns of attacker-versus-defender on the battlefield. One possible approach is discussed in [18].

Other generalizations of the Lanchester equations include:

- Partial differential equations to include maneuver; primarily work done by Protopopescu at the Oak Ridge National Laboratory
- Fuzzy differential equations to allow for imprecise information; see Dockery, [18]
- Stochastic differential equations to describe attrition processes under uncertainty

One can also speculate that there might be a way to generalize the Lanchester equations to include some kind of an internal *aesthetic*. That is to say, to generalize the description of the individual combatants to include an internal structure and mechanism with which they can adaptively respond to an external environment. See, for example, N. Smith's "Calculus of ethics," [58].

Nonlinear dynamics and chaos in arms-race models

G. Mayer-Kress [39] has written many papers on nonlinear dynamics and chaos in arms-race models and has suggested approaches to socio-political issues. His approach is to analyze computational models of international security problems using nonlinear, stochastic dynamical systems with both discrete and continuous time evolution. Many of Mayer-Kress' arms-race models are based on models of population dynamics first introduced by L. F. Richardson after World War I [51].

Mayer-Kress finds that, for certain ranges of values of control parameters, some of these models exhibit deterministic chaos. In one generalization of a discrete version of Richardson's equations that models the competition among three nations, for example, Mayer-Kress finds that the two weaker nations will form an alliance against the stronger nation until the balance of power shifts. The alliance formation factor and economical constraints induce nonlinearities into the model that result in multiple stable solutions, bifurcations between fixed point solutions and time-dependent attractors. He has also identified parameter domains for which the attractors are chaotic.

Tier IV: Description of the Complexity of Combat

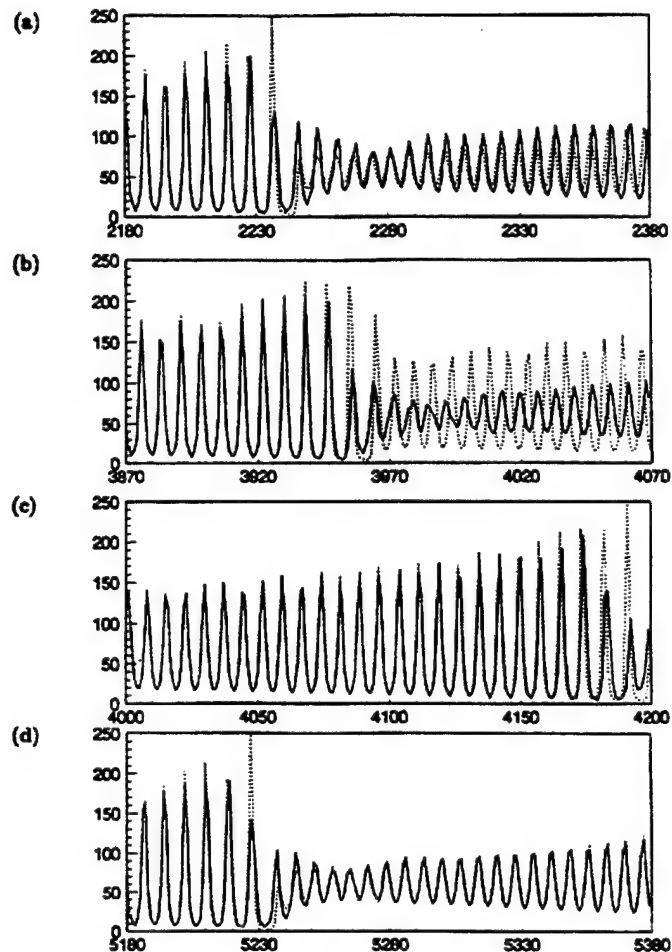
"The aim of science is not things themselves, as the dogmatists in their simplicity imagine, but the relations among things; outside these relations there is no reality knowable." – H. Poincare

Tier-IV consists of using the tools and methodologies of complex systems theory to describe and help look for patterns of real-world combat. It is the level on which complexity theory is effectively presented with a "candidate complex system" to study – that system being land combat – and given the opportunity to use its full arsenal of tools to explore in earnest the viability of this candidacy. Thus, this tier asks such basic questions as "What really happens on a battlefield?," "What kinds of complex systems theory inspired measures are appropriate to describe combat?," "Are there any embedded patterns, either in historical data or newly acquired data using sets of new measures, from which one can make short-term predictions of behavior?"

This tier consists of three sub-tiers of applicability:

- **Sub-Tier 1:** *Short-term predictability*, the objective of which is to exploit techniques such as attractor reconstruction to make short-term predictions about the progress of a battle or series of battles. Note that this does not require knowing the underlying rules governing the behavior of combat and/or having a working model
- **Sub-Tier 2:** *Confirmation of chaos from historical evidence*, the objective of which is to look for characteristic signs of underlying deterministic chaotic behavior in historical combat records. Some work has already been done in this area, most notably by Tagarev, *et. al.* [61], but much more remains.
- **Sub-Tier 3:** *Development of measures appropriate for describing combat from a complex systems theoretic point of view*. This sub-tier includes using such measures as Lyapunov exponents, power spectra, information dimension, and so on to redefine traditional data-collection requirements and measures-of-effectiveness of combat forces.

Figure 6. Four continuations of a chaotic time series using the embedding technique; solid lines represent predicted values, dashed lines represent the actual data



Attractor Reconstruction from Time-Series Data

Time-series analysis deals with the reconstruction of any underlying attractors, or regularities, of a system from experimental data describing a system's behavior (see Part I, pages 57-59). Techniques developed from the study of nonlinear dynamical systems and complex systems theory provide powerful tools whereby information about any underlying regularities and patterns in data can often be uncovered. Moreover, these techniques do not require knowledge of the actual underlying dynamics; the dynamics can be approximated directly from the data. These techniques provide, among other things, the ability to make short- (and sometimes long-) term predictions of trends in a system's behavior, even in systems that are chaotic.

Figure 6 shows an example of the kind of predictions that are possible with a popular technique called *time-delayed embedding*. Given 1000 data points (not shown) of the chaotic fluctuations in a far-infrared laser (approximately described by three coupled nonlinear ordinary differential equations) from which to learn the underlying system's dynamics, Sauer¹⁴ uses a modified embedding technique to predict the continuation of the time series for 200 additional time steps. Figure 6 (a-d) shows four continuations of length 200, each with a different initial point. In each of the plots, the solid curve represents the predicted continuation, and the dashed curve represents the true continuation.

Fractals and Combat


There are several suggestive fractal geometric aspects to land combat. For example, deployed forces are often assembled in a self-similar, or fractal, manner and are organized in a manifestly self-similar fashion: fire teams "look like" squads, which "look like" platoons, which "look like" companies, and so on. The tactics that are appropriate to each of these levels likewise shows the same nested mirroring. While a battalion is engaged in a frontal attack, for example, one of the companies could be conducting a supporting flanking attack that itself consists of two platoons engaged in a smaller-scale version of a frontal attack against two other platoons.

The FEBA (Forward Edge of the Battle Line) can also be characterized as a fractal, with greater and greater levels of detail emerging as the resolution is made finer. Woodcock and Dockery [18], for example, have plotted the FEBA length (in miles) versus the step size used in measuring the length using historical data from the German Summer Offensive of 1941 into Russia during World War II. Figure 7 shows three front line traces performed for three selected dates taken from the beginning, middle and end of the offensive. In each case, the log-log plot shows a close fit to the power-law fit characteristic of fractals: $L(h) \propto h^{1-D}$, where h is the measurement step size, L is the FEBA length and D is the fractal dimension.

¹⁴ Tim Sauer, "Time series prediction by using delay coordinate embedding," pages 175-193 in *Nonlinear Modeling and Forecasting*, edited by M. Casdagli, and S. Eubank, Addison-Wesley, 1992.

What this does or does not tell us about combat in general is an open question. The largest fractal dimension belongs to the FEBA trace recorded for the most active part of the campaign. Is this suggestive of something fundamental, or – because fractal dimension depends sensitively on the degree to which a line deviates from "straightness" – is it purely a consequence of the very convoluted nature of this particular trace? As a crude measure of the "complexity" of the evolving FEBA, the fractal dimension might be used to give a feel for the efficacy of a particular advance. Woodcock and Dockery suggest that the most immediate application of the fractal FEBA is to modeling, since it offers the possibility to generate a FEBA directly without detailed modeling. Moreover, a comparison between the daily changes in the fractal dimension of the FEBA calculated from an actual campaign and a computer model can be used to calibrate the model.

Figure 7. Power law scaling for FEBA length of German summer offensive of 1941 into Russia¹⁵



More generally, fractals dimensions of a variety of combat related systems (more examples are given below) can be used to quantify both the relevance of large-scaled events to the overall combat process and the subtle interrelationship that exists between small-scale events and large-scale outcomes.

Evidence of Chaos in War From Historical Data?


Tagarev, *et. al.* [61] provide extensive historical evidence of chaos in tactical, operational and strategic levels of military activity. Tagarev, *et. al.* examine (1) US fixed-wing aircraft losses during

¹⁵ Reference [18], page 321.

the Vietnam war, (2) US Army casualties in western Europe during World War II, and (3) historical trends in US defense spending.


As an example of chaos on the strategic level, consider figure 8, which shows a time-series plot of US aircraft losses in Vietnam. Is this just pseudo-random statistical fluctuation, or is it due in part to an underlying deterministic chaotic process?

Figure 8. Time dependence of US aircraft losses in Vietnam



One way to test for deterministic chaos is to look for the convergence of the estimated correlation dimension with increasing embedding dimension, shown in figure 9.

Figure 9. Estimated fractal dimension for weekly US aircraft losses in Vietnam (as a whole and in the air)



Recall that, loosely speaking, the fractal dimension of a set specifies the minimum number of variables that are needed to specify the set.¹⁶ Recall also that the embedding dimension is the

¹⁶ See Part I [28], pages 50-54.

dimension of the space in which the set of points making up the original time series (figure 8) is embedded.¹⁷ One essentially constructs d-dimensional data vectors from d points spaced equally in time and determines the correlation dimension of this d-dimensional point set. Since the set of data points making up the time-series consists of 443 points, the data points *a-priori* represent a 443-dimensional space. If the original data consisted of truly random points, then, as the embedding dimension is increased the calculated correlation dimension should also increase proportionately. The fact that the plot of fractal dimension versus embedding dimension seems to be converging as the embedding dimension increases suggests strongly that despite appearances, the irregular appearing time-series data shown in figure 8 is not random but is due to a deterministic chaotic process.

Evidence of Self-Organized Criticality From Historical Data?

Recall that self-organized criticality is the idea that dynamical systems with many degrees of freedom naturally self-organize into a critical state in which the same events that brought that critical state into being can occur in all sizes, with the sizes being distributed according to a power-law.¹⁸ Introduced in 1988, SOC is arguably the *only* existing holistic mathematical theory of self-organization in complex systems, describing the behavior of many real systems in physics, biology and economics. It is also a universal theory in that it predicts that the global properties of complex systems are independent of the microscopic details of their structure, and is therefore consistent with the "the whole is greater than the sum of its parts" approach to complex systems.

Combat Casualties

Is war, as suggested by Bak and Chen,¹⁹ perhaps a self-organized critical system? A simple way to test for self-organized criticality is to look for the appearance of any characteristic power-law distributions in a system's properties. Richardson [51] and Dockery and Woodcock [18] have examined historical land combat attrition data and have both reported the characteristic linear power-law scaling expected of self-organized critical systems. Richardson examined the relationship between the frequency of "deadly quarrels" versus fatalities per deadly quarrel


¹⁷ See Part I [28], page 49.

¹⁸ See Part I [28], pages 101-107.

¹⁹ P. Bak and K. Chen, "Self-Organized Criticality," *Scientific American*, Volume 26, January 1991, 46-53.

using data from wars ranging from 1820 to 1945. Dockery and Woodcock used casualty data for military operations on the western front after Normandy in World War II and found that the log of the number of battles with casualties greater than a given number C also scales linearly with $\log(C)$; see figure 10.

Figure 10. Analysis of WWII casualty data on the western front after Normandy (Dockery and Woodcock, [18])




The paucity of historical data, however, coupled with the still controversial notions of self-organized criticality itself, makes it difficult to say whether these suggestive findings are indeed pointing to something deep that underlies all combat or are merely "interesting" but capture little real substance. Even if the results quoted above do capture something fundamental, they apply only to a set of many battles. The problems of determining whether, or to what extent, a power-law scaling applies to an individual battle or to a small series of battles, and – perhaps most importantly – what tactically useful information can be derived from the fact that power-law scaling exists at all, remain open.

Message Traffic

Woodcock and Dockery [18] also examine message traffic delays using data collected as part of a military exercise; see figure 11. They find that the number of messages that arrive after a given time delay again follows a linear power-law scaling expected of self-organized critical systems. The authors comment that

"The conditional expectation of further delay after waiting a given period actually increases in proportion to the time already waited. Compare this property with the more usual assumption of the Poisson distribution that the delay is independent of the time already waited. The fractal distribution is thus much more in accord with the maxim that 'the worse things get, the more worse they can get.' For the operational commander, the consequence of the hyperbolic fit is that self-initiated action is probably called for after a suitable delay. For the message traffic system designer the implication of the power law fit is to make messages, which are long delayed, candidate for deletion."

Figure 11. Log-log plot of message delay in 2.5 minute time bins



Use of Complex Systems Inspired Measures to Describe Combat

Part I of this report [28] provides a detailed discussion of several useful qualitative and quantitative characterizations of chaos. Qualitative characterizations include time-plots of the behavior of pertinent variables, *Poincare plots*, *autocorrelation functions* and *power spectra*. Quantitative characterizations include *Lyapunov exponents*, *generalized fractal dimensions* (including fractal, correlation and information dimensions), and *Kolmogorov-Sinai entropy*. A recently introduced idea is to use casualty-based entropy as a predictor of combat.

Casualty-Based Entropy

Carvalho-Rodrigues [13] has recently suggested using entropy, as computed from casualty reports, as a predictor of combat


outcomes. Whether or not combat can be described as a complex adaptive system, it may still be possible to describe it as a dissipative dynamical system (see Part I [28], page 28). As such, it is not unreasonable to expect entropy, and/or entropy production, to act as a predictor of combat evolution. Carvalho-Rodrigues defines his casualty-based entropy E by

$$E_i = \frac{C_i}{N_i} \log \frac{1}{C_i/N_i} ,$$

where C_i represents the casualty count, in absolute numbers) and N_i represents the force strength of the i th adversary (either red or blue). It is understood that both C_i and N_i can be functions of time.

Figure 12 shows a plot of the functional form $E(x) = x \log (1/x)$, where $x = C_i / N_i$. Notice that the curve is asymmetrical and has a peak at about 0.37. One could interpret this to mean that once C_i / N_i goes beyond the peak, "it is as if the combat capability of the system ... declines, signifying disintegration of the system itself."²⁰

Figure 12. A plot of entropy $E(x) = x \log (1/x)$




Woodcock and Dockery [18] provide strong evidence that casualty-based entropy is a useful predictor of combat. They base this on analysis on both time-independent and time-dependent combat data derived from detailed historical descriptions of 601 battles from circa 1600 to 1970, exercise training-data obtained from the National Training Center and historical records of the West-Wall campaign in World-War II and Inchon campaign during the Korean war.

They find that plots of E_a (attacker entropy) versus E_d (defender entropy) are particularly useful for illustrating the overall combat process (see figure 13):

- *Region I*: a low entropy region corresponding to low casualties and ambiguous outcomes. Initial phases of a battle pass through this region, with the eventual success or failure for a given side depending on the details of the trajectory in this entropic space
- *Region II*: a region of high entropy for the defender and low entropy for the attacker indicates the attacker wins
- *Region III*: a region of ambiguous outcomes, like region I, region III represents high attrition with outcomes depending on the direction of the trajectory. (Woodcock and Dockery indicate that only simulated combat appears able to reach this region.²¹)
- *Region IV*: an analogue of region II, where the entropy roles are reversed and the defender wins.

Figure 13. Regions of casualty-based entropy phase space [18]



Woodcock and Dockery further suggest that the measurement and display of coupled casualty and reinforcement rates may be a first step towards quantifying the *battle tempo*. "The tempo is then seen to characterize, not the physical rate of advance (the usual

²¹

Reference [18], page 223.

connection), but rather the rate of structural breakdown of the fighting force."²²

We note, in closing, that Carvalho-Rodrigues's definition of a casualty-based entropy is but one possible definition. One could alternatively use generalizations of the Renyi-entropy, Kolmogorov-Sinai entropy, or topological entropy, among many other definitions. Despite the seeming simplicity of the basic idea, there is strong evidence to suggest that entropy will play a fundamental role in understanding the underlying dynamical processes of war.

Use of Relativistic Information to Describe Command and Control Processes

Relativistic information theory is a concept introduced by Jumarie [29] and has been suggested as a possible formalism for describing certain aspects of military command and control processes by Woodcock and Dockery [18]. Relativistic information may prove to be particularly useful for gaining insight into the interplay among combat, command and control and information.

Generalized entropy is an entropy that is endowed with four components, so that it is equivalent to a four-vector and may be transformed by a Lorentz transformation (as in relativity theory). These four components consists of:

1. *external entropy* of the environment (H_o), which can be associated with operational and intelligence information
2. *internal entropy* of the system (H_i), which can be associated with the readiness of forces information
3. system goals, which can be equated with missing (and planning) information
4. *internal transformation potential*, which measures the efficiency of the system's internal information transformation; this can be associated with a measure of the command and control capability information

An additional factor, called *organizability*, plays the role of "velocity." Woodcock and Dockery show that it is possible to use relativistic information theory to compare the relative command and control system response of two command structures to the world around them. The quantity of interest is dH_i/dH_o , or the rate of change of the internal information environment with respect to changes in the surrounding environment.

"Using relativistic information theory it is possible to compare the relative command and control system response of two commanders to the world around them. Their relative perceptions of the change about them is theoretically quantified by relativistic information theory. Because the theory measures changes with respect to the environmental change, we can argue that self-organization is a requirement for a military force. If the internal structure cannot cope with the change in the environment that structure must itself change. The goal of combat must paradoxically be to create a self-organizing structure which nonetheless ensures the destruction of the foes' internal structure."²³

Tier V: Combat Technology Enhancement

Tier V consists of applying complex systems theory tools to enhance existing combat technologies. This "workhorse" tier is concerned with using specific methods to improve, or provide better methods for applying, specific key technologies. Examples include using computer viruses (a form of "artificial life") as computer countermeasure agents, applying iterated function systems (i.e. fractals) to image compression for data dissemination, using cellular automata for cryptography, using genetic algorithms for intelligent manufacturing, using synchronized chaotic circuits to enhance IFF capability, and "fire-ant" technology.

Computer viruses ("computer counter-measures")

A computer virus can be thought of as an autonomous agent. It is a computer program that tries to fulfill a goal or set of goals without the intervention of a human operator. Typically, of course, viruses have rather simple and sinister goals of tampering with the normal operation of a computer system and/or computer network and then reproducing in order to spread copies of themselves to other computers. Computer viruses are particularly interesting to artificial life researchers because they share many of the properties of biological viruses.

From a military standpoint, computer viruses can be used in two ways: (1) as computer countermeasure agents to infiltrate enemy systems, or (2) as constructive "cyberspace allies" that, for example, can be programmed to maintain the integrity of large databases.

Fractal Image Compression

A powerful technique for image compression that is based on fractals – called *Iterated Function Systems* (IFS) – has been developed by Barnsley and his co-workers at the Georgia Institute of Technology [4]-[5]. To appreciate the need for compressing images, consider a typical grey-scale intelligence photograph that need to be disseminated to interested parties. Suppose there are 256 shades of grey and that the image must be scanned and converted into a 1024-by-1024 pixel digitized image. The resulting image can be recorded using a binary string of 1024-by-1024-by-8 binary bits of information. Thus, without

compression, one must use close to 10 million bytes of memory to store the image. Image compression involves reducing the required number of bytes to store an image, and can be either "lossless," so that the original image can be recovered exactly, and "lossy," so that only an approximate version of the image can be retrieved.

There are, of course, tradeoffs involved among how well the original image can be recovered, what the maximum possible compression rate is, and how fast the actual compression algorithm can be run. Generally speaking, the greater the desired compression, the more CPU-time is required and the greater the risk of some compression loss. Conventional lossy compression schemes, such as Discrete Sin and Cos Transforms, can achieve compression ratios ranging from 2:1 to 10:1, depending on the image. In comparison, while IFS is generally lossy (so that an original image cannot generally be recovered from the compressed image exactly), it is able to achieve extremely high compression ratios, approaching 50:1, 300:1 or better. *Microsoft* has, in fact, licensed use of this technology to compress images found on its CD-ROM encyclopedia *Encarta*.

The basic idea of IFS is simple to state, though often time-consuming to apply without special hardware. An image is compressed by exploiting the innate self-similarity, or redundancy, contained in an image. Recall that fractals are, loosely speaking, objects that consist of an endless succession of smaller versions of themselves, at all levels.

IFS uses *affine transformations* to build a "collage" of an image using smaller copies of the image. An affine transformation is an operation on a set of points that distorts that set by scaling, shifting, rotating and/or skewing. The IFS process involves finding smaller, distorted copies of an image and putting them together in a "collage" that approximately reproduces the original image. Each distorted copy of the image represents a different affine transformation. Once a collage is formed, the original image can be thrown away. The original image can be recovered by applying the appropriate set of affine transformations to some starting "seed" coordinate and iterating. The trajectory of the points will converge onto an attractor that defines the image. For example, figure 14 shows a "recovered" image of a fern using the affine transformation defined by $(x,y) \rightarrow (s_1 x \cos r_1 - s_2 y \sin r_2 + t_1, s_1 x \sin r_1 + s_2 y \cos r_2 + t_2)$, where the parameters r_1, r_2, s_1, s_2, t_1 and t_2 are shown in table 7.

Table 7. Rotation (r_1 , r_2), scaling (s_1 , s_2) and translation (t_1 , t_2) parameters defining the IFS affine transformation of a fern

Affine Transformation	r_1	r_2	s_1	s_2	t_1	t_2
1	0	0	0	0.16	0	0
2	-2.5	-2.5	0.85	0.85	0	1.6
3	49	49	0.3	0.34	0	1.6
4	120	-50	0.3	0.37	0	0.44

Consider what this fern represents. A high resolution digitized image of the original grey-scale image takes up more than a megabyte of memory. Conventional compression schemes might reduce this by a factor of five. But IFS has reduced the image to essentially 28 parameters.

Figure 14. An IFS encoded fern using the affine transformation defined in table 7



Of course, not all objects in nature have the manifest self-similarity of a fern. The trick is to find the right group of affine transformations for generating a given image. That a set of affine transformations can be found in general, even for objects that *do not* exhibit a manifest self-similarity is due to a theorem called the *Collage Theorem*. The collage theorem asserts that given a target image S and a measure of closeness ϵ , a set of IFS affine transformations f can be found such that the "distance" (appropriately measured) between S and $f(S)$ (which is the

union of scaled copies of S , called the "collage" of S) is less than ϵ [5].

Applications of IFS compression are wide-ranging and far-reaching. They include more efficient transmission of fax, still imagery and video, more efficient computer data storage and retrieval, and image recognition. There are of course many more details to IFS than there is room to discuss in this paper. An excellent reference is Barnsley's book [5]. His company, *Iterated Functions*, also has an extensive collection of articles and software demos on the WWW at <http://www.iterated.com>.

Cryptography

In succinct terms, an ideal cryptographic encryption scheme is an operation on a message that renders that message completely meaningless to anyone who does not possess a decryption key, and, at the same time, preserves and reveals the original message exactly to anyone who possess the key. Ideally, the operation is encrypts quickly and decrypts quickly.

All practical cryptographic schemes, of course, are less than ideal, typically because their encryption schemes are less than foolproof (Denning, [16]). Most depend on the presumed computational difficulty of factoring large prime numbers. The effective measure of worth of any cryptographic scheme remains the test of time: the longer a given system is in widespread open use by trained intelligent cryptanalysts without being "broken," the better the system.

It would take us too far afield of the main subject of this report to go into any great detail about cryptanalysis. We will only briefly mention some attempts that have been to develop cryptosystems based on nonlinear dynamical system theory and cellular automata (See Part I [28], pages 81-88). We follow mainly Gutowitz ([24]).

The basic idea is to use a nonlinear dynamical system that is known to exhibit deterministic chaos and use it to evolve some initial starting point to some future state. After a certain time has been allowed to elapse, the initial state – which is defined to be the public key – is effectively "forgotten." Because the dynamics is assumed to be strictly deterministic, however, the same initial state always leads to the same final state after a specified number of time steps. Users can thus send messages to one another

encoded in some way using some part of the evolved trajectory traced out by the secret initial state. Anyone who does not possess the initial state will not be able to reproduce the trajectory and thus will not be able to decipher the message.

Bianco and Reed [9] have patented an encryption scheme using the logistic equation as the underlying dynamical system. A drawback to this scheme, however, is that the sequences generated by the logistic map are not truly random, so that an appropriate statistical analysis could identify embedded patterns that could then be exploited to decipher a message. Wolfram [68] suggests a discrete dynamical system version of the basic idea that uses the iteration of a cellular automaton to generate a bit string. The cellular automaton chosen (shown in figure 20 on page 84 in Part I [28]) is known as rule 30. What is interesting about this rule is that the temporal sequence of vertical values of its evolving space-time pattern has been shown to satisfy all known tests of randomness. As for the case of a continuous dynamical system, the secret key is the initial state of the cellular automaton system, and a message can be encrypted and decrypted by combining it with the temporal sequences of a given length generated by the rule. Gutowitz [24] has also introduced a much more powerful and sophisticated algorithm based on a cellular automaton model.


Tier VI: Combat Aids

Tier VI consists of using the tools of complex systems theory to enhance real-world operations. Examples include using genetic algorithms to "evolve" strategy and/or tactics, developing tactical picture agents to adaptively filter and integrate information in real-time, and developing autonomous robotic devices to act as sentries and data collectors.

As has been discussed elsewhere (see Part I, pages 93-101), genetic algorithms are powerful heuristic tools for finding near-optimal solutions for general combinatorial optimization search problems. One obvious application of genetic algorithms that has found a comfortable home in the artificial life research community, involves their use as sources of the "adaptive intelligence" of adaptive autonomous agents in an agent-based simulation. A related application that is of particular interest to the military strategist, theorist and/or battlefield commander, is that of *direct strategy and/or tactics development*.

Figure 15 shows a schematic representation of what might be called a "strategy landscape." The strategy landscape represents the space of all possible global strategies that can be followed in a given context or scenario. Generally speaking, a genetic-algorithm-based tactics- or strategy- "optimizer" consists of an evolutionary search of this landscape for high-pay-off strategies using whatever local information is available to individual combatants. The shape of the landscape is determined by the fitness measure that is assigned to various tactics and/or strategies. It also changes dynamically in time, as it responds to the actual search path that is being traversed.

Figure 15. Schematic representation of a strategy landscape



Using Genetic Algorithms to Evolve Tank Strategies

Carter, Dean, Pollatz, *et. al*, [12] suggest using genetic algorithms to "evolve" strategies for the battlefield. While their ultimate goal is to develop a complete architecture for intelligent virtual agents that employ multiple learning strategies, their initial testbed consists of evolving reasoning-strategies for what they call *smart tanks*. While this testbed is deliberately designed to be as simple as possible, because it involves many of the key elements that make up more realistic models, it is of considerable pedagogical value. For this reason, we discuss the smart-tank testbed briefly below.


Smart tanks live on a simple two-dimensional "battlefield" containing a randomly placed "black-hole" (see figure 16). The black-hole represents a lethal area of the battlefield that annihilates any smart tank that encounters it. Smart tanks, generated on one side of the battlefield, must cross over to the other side without encountering the black-hole if they are to be successful and "live." A tank's route is determined by its *genotype* (see below).

A smart tank is an artificial organism that consists of three basic components: (1) *memory*, which is a record of the decisions and fates of three previous smart tanks that have successfully crossed the battlefield, (2) *reasoning*, which is the internal mechanism by which a tank selects one of several viable strategies, and (3) *instinct*, which is a basic inference engine common to all tanks, and is included to simulate the basic behaviors that may occur in life threatening or otherwise critical situations. How much weight a given tank assigns to each component – that is, what overall reasoning strategy it chooses to follow – depends on its genetic predisposition.

Smart-tank strategies "evolve" in the following way. First, a population of smart tanks is created. One tank is selected out of this population and crosses the battlefield. It looks ahead a certain number of discrete bins into which the battlefield is decomposed, and gathers information on the destination bin. The exact number of bins that it looks ahead at depends on the tank and its current reasoning strategy. The tank then selects a reasoning strategy upon which to base its next move. After arriving at the destination square this same process repeats itself. If, at any time, the tank encounters a black-hole, it dies and its genetic structure is lost (though a record of it is maintained in a

global case-base). If the tank survives – i.e. never encounters – the black-hole, its genes are saved in a "winner's circle." A genetic algorithm then combines the genes in the winner's circle to form a new population of tanks.

Figure 16. Schematic representation of a genetic-algorithm based "smart tank"



Carter, *et. al.*'s initial testbed was designed to use genetic algorithms as a mechanism for creating coordination schemes for three specific learning strategies:

- **Case-based Reasoning.** This strategy consists of using whatever reasoning was the best approach for an identical, or almost identical, situation encountered in the past. Carter, *et. al.*'s actual implementation involved tacking a history of up to 256 previous tank traversals of a 10-by-10 bin battlefield, with a maximum of three histories being accessible by any one tank. The case-based reasoning strategy compares the three histories and selects the closest fit. The suggested move is assigned a confidence level commensurate with the closeness of fit.
- **Rule-based Reasoning.** This strategy consists of using the current information about the local surroundings and determining what rule, out of the current rule set, is best applicable. Rules are of the form "go to bin nearest goal," or "go to bin requiring least amount of energy to get to," and so on. The rule-base also takes into account how well rules have performed in the past.

- **Instinct.** This strategy consists of using one of a set of generic responses. Moves are based on prior determined moves for a given bit pattern in a tank's chromosome. A typical move might consist of going to the adjacent bin that exerts the least amount of pull toward the black-hole (a black-hole's "attraction" for tanks diminishes with distance from black-hole, but its effects are felt in many surrounding bins).

No one of these individual reasoning strategies, of course, works equally as well in all situations. Indeed, one of the main reasons for developing this testbed example was to explore the efficacy of various options and to allow the genetic algorithm to suggest the right "mix" of strategies.

Smart-Tank's Chromosome

A smart tank's actual chromosome consists of 35 genes. We should immediately emphasize that as is true of almost all other features of this simple testbed, there is nothing sacrosanct about having 35 genes. One could choose to have a greater or lesser number of genes, and to interpret the genes in a different manner from that outlined below. The testbed is here presented in detail for illustrative purposes only, and to suggest only one of many equally as valid approaches that could be used to design a genetic algorithm scheme for evolving strategies for the battlefield.

The functions of a smart tank's 35 binary-valued genes (i.e. each gene takes on either the value 0 or 1) are broken down as follows:

- **Bits 1 - 8:** pointer to the *first* of 3 accessible histories of previous tank traversals (out of a maximum of 256 stored histories; see "Case-based reasoning" above)
- **Bits 9 - 16:** pointer to the *second* of 3 accessible histories of previous tank traversals
- **Bits 17 - 24:** pointer to the *third* of 3 accessible histories of previous tank traversals
- **Bits 25 - 29:** *tank characteristics* (2 bits for sensitivity, 1 bit for speed, and 2 bits for type and range of look ahead)

- **Bits 30 - 32:** *type of reasoning* (the 3 bits specify 9 3-bit patterns ranging from 100% rule-based for pattern 111 to 50% rule-based and 50% case-based for pattern 100 to 33% rule-based, 33% case-based and 34 % instincts for pattern 011).
- **Bits 33 - 35:** generic instincts (the 3 bits again specifying 9 possible 3-bit patterns that range from "tank heads toward the goal regardless of forces in the environment" for pattern 000 to "tank moves to square in front of it that has the least force in it" for pattern 010, and so on)

As mentioned above, a greater or fewer number of genes could have been chosen, and all or some of their interpretations altered. What it is hoped the reader will take away from this description of a simple testbed is the general approach to building a genetic algorithm based "strategy evolver."

Tactical Decision Aids

Virr, Fairley and Yates [65] have suggested using a genetic algorithm as an integral component of what they call an Automated Decision Support System (ADSS).

There are three main phases to any decision making process:

1. *Data Fusion:* wherein all of the available information is assembled to form a tactical picture of a given situation
2. *Situational Assessment:* wherein various pertinent aspects of the tactical picture are appraised
3. *Decision:* wherein the appropriate action, or set of actions, is actually selected

The tactical plan, from which a specific set of condition-contingent actions is selected, can be expressed, in its simplest form, as a sequence of IF-THEN rules of the form

.....
Rule R_{n-1}
Rule R_n :
IF(condition₁ AND condition₂ AND ... AND condition_N)

THEN(action₁ AND action₂ AND ... AND action_M)

Rule R_{n+1}

.....

Here, condition n_i refers to the i^{th} piece of information assimilated during the data fusion phase (such as speed of target, bearing, etc.) and action_j refers to the j^{th} action that is to be taken (such as turn toward target, engage target, and so on).

A Plan P, consisting of a set of rules $\{R_1, \dots, R_p\}$, forms the rule-base of a Knowledge Based System (KBS) for a given tactical situation. Once specified, it can be used for training purposes, to model tactical decisions made by an adversary in simulated combat and/or as a real-time tactical decision aid on the battlefield. The problem, of course, is how to construct P.

Traditionally, P has been constructed by a knowledge engineer; that is, someone who painstakingly elicits from experts the facts and heuristics used by those experts to solve a certain set of problems. Of course, there are two obvious problems with this traditional approach: (1) an expert may not always be able to successfully articulate all of the relevant knowledge required for solving a problem, and (2) there may be enough of a mismatch between the concepts and vocabulary as used by an expert and a knowledge-engineer that though an expert may correctly articulate the relevant knowledge, the knowledge engineer is unable to render that knowledge meaningful within the IF-THEN rule structure of the plan.

One way of circumventing both of these problems – called Machine Learning (ML) – is to have the KBS "discover" the required knowledge, and thereby construct the plan P, by itself. Although there are many different techniques that all go under the rubric of ML, they all fall into one of four major categories:

- **Analytic Learning.** Analytic learning systems require a thorough understanding of the general underlying problem type and must have available a large number of problem-solution exemplars. The technique relies on adapting solutions to problems that it identifies as being "close to" known solutions to known problems.
- **Inductive Learning.** Inductive learning requires an external "teacher" to produce problem samples. The teacher grades the system's attempts to use its stored knowledge to try to

solve each problem in turn. The teacher's grade is then used to update the system's knowledge.

- **Neural Network Learning.** The neural net (also called the *Connectionist*) approach consists of applying a learning algorithm (such as *back-propagation*) to adjust a set of internal weights in order to minimize the "distance" between calculated and desired solutions to selected problems. Given a set of training problem-solution exemplars, the learning algorithm produces a network that, in time, is able to correctly recognize the pattern implicit in all input (i.e. problem) and output (i.e. solution) pairs.
- **Genetic Algorithm (or *Selectionist*) Learning.** Selectionist learning systems exploit the learning capability of a genetic algorithm to "evolve" an appropriate knowledge base. Recall that genetic algorithms are a class of heuristic search methods and computational models of adaptation and evolution that mimic and exploit the genetic dynamics underlying natural selection.

Given the basic differences among these four approaches, it is clear that, in general, not all approaches can be expected to be equally appropriate for solving a given kind of problem. Depending on the problem, each approach offers certain unique advantages and disadvantages.

When it comes to the general problem of tactical decision making, a strong case can be made that selectionist learning techniques are the most appropriate. First, there is no complete "domain theory" describing all possible conflicts and scenarios on which to base a general strategy of conflict resolution. This makes it hard to use an analytical learning technique. Second, whatever real-world expertise there is to assist in building a KBS must, of necessity, be both incomplete (because only a small fraction of all possible scenarios can be experienced) and imprecise (because all human experience is fundamentally subjective). Thus, both inductive and connectionist learning techniques, both of which depend critically on having sets of carefully pre-constructed scenario-plan exemplars available for learning, would be difficult to use for this problem. Finally, and most importantly, any tactical plan must be able to continually adapt to changing, and often unanticipated, facts and scenarios. Genetic algorithms, of course, are designed to deal with precisely

this kind of open-ended and "changing" problem, as they are particularly adept at discovering new rules.

Now, how specifically can genetic algorithms be used to generate new tactics? The answer depends on how genetic algorithms are incorporated within the parts of a larger class of *Classifier System*.

Classifier Systems

Classifier systems were introduced by John Holland as an attempt to apply genetic algorithms to cognitive tasks. A classifier system typically consists of (1) a set of *detectors* (or input devices) that provide information to the system about the state of the external environment, (2) a set of *effectors* (or output devices) that transmit the classifier's conclusions to the external environment, (3) a set of *rules* (or classifiers), consisting of a condition and action, and (4) a list of *messages*.

Rules are the actual classifiers, and are grouped together to form the classifier's rule-base. Associated with each classifier is a *classifier-weight*, representing the degree of usefulness of that particular classifier in a given environment. Messages constitute the classifier system's basic means of exchanging information, both internally and at the interface between classifier system and external world.

Although the operation of a real classifier system can be quite complex, their basic operation consists of the following general steps: Information from the world model is first communicated to the classifier at the input interface. The classifier combines this information with rules stored in its rule-base to select an appropriate action, which is, in turn, effected at the output interface, updating the world model. Learning takes place via *credit assignment*, wherein rules are judged "good" or "bad" in order to teach the system what actions are appropriate in what contexts. The genetic algorithm comes in as the part of the classifier system responsible for deciding how "old" rules in the rule-base are replaced by "new" rules.

How can Genetic Algorithms be Used?

Recall that genetic algorithms process a population of "solution-organisms" according to their relative "fitness" (i.e. a figure of merit ostensibly measuring an organism's ability to solve a given problem) so that, over time, as the population evolves, there is an increasing likelihood that some members of the

population are able to "solve" the given problem well (or well enough). In the present context, the "problem" is to find new and/or improved rules for a tactical decision knowledge base. The meta-problem, from the point of view of the genetic algorithm, is how to go about ascribing a "fitness" to members of the rule-population. Without a fitness, of course, there is no way for the population to evolve.

Virr, et. al. [65] suggest four ways in which a genetic algorithm can be grafted into a classifier system to effectively breed rules:

1. *Apply Genetic Algorithm at the Rule Level.* Suppose we take a "population" to consist of rules making up a particular plan. A genetic algorithm can then use the individual rule-strengths as fitnesses guiding their evolution. The major drawback to this approach is that since all of the rules are independent, the genetic algorithm degenerates into a search for a "super-rule" that deals with all situations (which, for typical real problems, does not exist). There are ways, however, of inducing rules to cooperatively link with one another, partially circumventing the drawback to this approach.
2. *Apply Genetic Algorithm at the Plan Level.* An alternative is to use a genetic algorithm on a population of plans rather than on a population of rules making up a given plan. Drawbacks to this approach include (1) using a single fitness measure (presumably derived from the fitnesses of the individual rules) to represent the efficacy of an entire plan, and (2) the need for an additional algorithm to generate new rules in the various plans.
3. *Apply Genetic Algorithm at the Sub-Plan Level.* Suppose a plan P is partitioned into q subsets, where $1 < q < R$, R is the number of rules in P, and the rules in each subset of the partition are related in some way. Then the rules within a each subset can be viewed as a population, and – since each rule has an associated rule-weight – the population can be subjected to a genetic algorithm. The efficacy of the partitioning scheme itself may also be amenable to a genetic algorithm.
4. *Apply Genetic Algorithm at Both the Rule and Plan Levels.* The fourth approach attempts to capitalize on the advantages of the first two approaches by carefully combining them.

The idea is to associate one classifier with each plan, and running the set of classifiers in parallel. The genetic algorithm is then applied both at the rule level of each classifier – during which time the plans are allowed to develop independently – and at the plan level, during times in which the operation of the individual classifier systems is periodically suspended.

Tactical Picture Agents

Anyone who has spent even a small amount of time "surfing" the World-Wide-Web for information can attest to how difficult it is to find useful information. To be sure, the WWW is filled with untold numbers of glossy pages overflowing with all kinds of information. A quick use of a web search-engine such as *Lycos*²⁴ or *AltaVista*²⁵ usually suffices to uncover some useful sites. But what happens when one needs to find some information about a particularly obscure subject area? And what happens when one begins relying on one's web connection for more and more of one's daily workload: e-mail, stock quotes, work scheduling, selection of books, movies, travel arrangements, video conferencing, and so on?

A powerful emerging idea that helps the human "web-surfer" deal with this increasing workload and that is based in part on the methodologies of autonomous agents and genetic algorithms, is that of *Intelligent Software Agents*.²⁶

Software agents are programs that essentially act as sophisticated personal assistants. They act as intermediaries between the interests of the user and the global information pool with which the user has traditionally dealt directly. Software agents engage the user in a cooperative process whereby the human operator inputs interests and preferences and the agent monitors events, performs tasks, and collects and collates useful information. Because software agents come endowed with an adaptive "intelligence," they become gradually more effective at their tasks as they begin learning the interests, habits and preferences of the user.

²⁴ <http://lycos.cs.cmu.edu/>.

²⁵ <http://altavista.digital.com/>.

²⁶ See, for example, the collection of articles in *Communications of the ACM*, Volume 37, No. 7, July 1994.

How does this relate to a military combat environment? Intelligent software agents can be used for *adaptive information filtering and integration* and as *Tactical Picture Agents*, scouring and ordering the amorphous flood of battlefield and intelligence data.

ONR Initiative

The basic idea is rapidly nearing the development stage. As an example of one application that is already in the embryonic stage of development for the navy, the Office of Naval Research (ONR) has recently released a public announcement soliciting submission of research proposals for a 5-year research initiative to develop intelligent tactical picture agents for naval decision-makers.²⁷

Naval commanders must have access to, and have immediate use of, the right information at the right time. They must assimilate and understand all of the relevant information concerning their own situation, including the disposition of the Red, Blue and White forces, geographical, oceanographic and meteorological characteristics of the surrounding vicinity, status of all weapon systems, and so on. The totality of this information is called the *Tactical Picture*.

At the present time, information is disseminated typically via either naval text messages, ship-to-ship communications, radar and sonar tracks, and so on. As technology improves, naval ships and other platforms will have access to a wider range of information, including immediate access to satellite images and weather reports, on-line intelligence analyses, and perhaps even direct connectivity to the world-wide-web or some similar globally connected network. It therefore becomes vital to develop intelligent software agent technologies that can automatically perform the data-mining and filtering functions necessary to make effective use of this potential explosion of information.

ONR's 5-year basic and exploratory research effort is designed to foster the development of a tactical picture agent technology that can eventually dramatically improve the on-board "tactical picture building" ability of all naval platforms. ONR has spelled out three basic objectives:

²⁷ An HTML-formatted copy may be obtained from the WWW URL address <http://jhuapl.edu/program/tpa/>. Other information can be obtained at <http://www.itd.nrl.navy.mil/ONR/aci/tpahome.html>.

1. Modeling of users and tasks so that intelligent software can decide what to search for and how to integrate search results.
2. Development of intelligent software agents that can locate and filter multi-media information appropriate for a particular task and user.
3. Develop methods of displaying information that are appropriate for a particular task and environment.

Exactly the same ideas apply to developing a tactical picture agent for land combat. Diverse forms of information must be assimilated, filtered, ordered and presented to the field commander.

Autonomous Robotic Devices

Adaptive agent technology can be used to develop autonomous robotic devices to act as sentries, to help in material transportation and hazardous material handling.

Tier VII: Synthetic Combat Environments

Tier VII consists of developing full system models for training purposes and/or for use as research laboratories. Examples include cellular-automata-based and multi-agent-based simulations of combat (along the lines of commercial games like *SimCity*), cognitive architecture driven simulations such as the one found in Carnegie-Mellon University's SOAR/IFOR project, and combat models based on the *Santa Fe Institute's* SWARM general-purpose complex adaptive modeling system.

Combat simulation using cellular automata


If one abstracts the essentials of what happens on a battlefield, ignoring the myriad layers of detail that are, of course, required for a complete description, one sees that much of the activity appears to involve the same kind of simple nearest-neighbor interactions that define cellular automata (see Part I, page 81). Woodcock, Cobb and Dockery [18] in fact show that highly elaborate patterns of military force-like behavior can be generated with a small set of cellular automaton-like rules.

In Woodcock, *et. al.*'s model, each combatant – or automaton – is endowed with a set of rules with which it can perform certain tasks. Rules are of four basic varieties:

- *Situation Assessment*, such as the determination of whether a given automaton is surrounded by friendly or enemy forces
- *Movement*, to define when and how a given automaton can move; certain kinds of movement can only be initiated by threshold and/or constraint criteria
- *Combat*, which governs the nature of the interaction between opposing force automata; a typical rule might be for one automaton to "aim fire" at another automaton located within some specified fight radius
- *Hierarchical Control*, in which a three-level command hierarchy is established; each lower-level echelon element keys on those in the next higher echelon on each time step of the evolution

These basic rules can then be augmented by additional rules to (1) simulate the impact that terrain barriers such as rivers and mountains have on the movement of military forces; (2) provide a capability for forces to respond to changing combat conditions (for example, a reallocation of firepower among three types of weapons: aimed firepower, area firepower and smart weapons firepower), and (3) replace entities lost through combat attrition. Figure 17 shows a schematic of three sample rules. A further extension involves relating notional features of battlefield geometry to the structure of real battlefields [18].

Figure 17. Three sample rules in Woodcock, *et. al.*'s CA combat model



Woodcock, *et. al.* stress that the goal of CA-based model of combat is *not* to codify a body of rules that comes as close as possible to the actual behavioral rules obeyed by real combatants; rather, the goal lies in "finding the simplest body of rules that both can generate nontrivial global combat-like phenomena and provide a new understanding of the combat process itself by extracting the maximum amount of behavioral complexity from the least complicated set of rules." [18] Additional details are discussed in chapters 3.1 and 3.2 of reference [18].

Agent-based simulations

For many obvious reasons, the most natural application of complexity theory to land warfare is to provide an *agent-based simulation of combat*. The basic idea is to model land combat as a co-evolving ecology of local-rule-based autonomous adaptive agents.


An Irreducible Semi-Autonomous Adaptive Combat Agent (ISAACA) represents a primitive combat unit (infantryman, tank,

transport vehicle, etc.) that is equipped with the following characteristics (see figure 18):

- a default local rule set specifying how to act in a generic environment (i.e. an embedded "doctrine")
- goals directing behavior ("mission")
- sensors generating an internal map of environment ("situational awareness")
- an internal adaptive mechanism to alter behavior and/or rules; adaptation is *genetic-algorithm*-based (see page 93) – each ISAACA effectively plays out a scenario using a genetically-encoded set of possible tactics; where fitness is the expected payoff with respect to some internal value system

An ISAACA collective, represented schematically in figure 19, consists of local and global commanders, each with their own command radii, and obeys an evolving C^2 hierarchy of rules. A global rule set determines combat attrition and reinforcement. Nonlinear feedback exists among combatants (measure \rightarrow countermeasure \rightarrow countercountermeasure \rightarrow ...) and between combatants and the environment.

Figure 18. Field-of-view of a single ISAACA



Note that this approach is similar in spirit to a *cellular automaton* (CA) model (see Part I, page 81) but augments the conventional CA framework in three ways: (1) evolution proceeds not

according to a fixed set of rules, but to a set of rules that adaptively evolves over time; (2) individual states of cells (or combatants) do not just respond to local information, but are capable of non-local information (via an embedded C^2 topology) and command hierarchy; and (3) global rule (i.e. command) strategies are evolved via a genetic algorithm (orders pumped down echelon are based on evolved strategies played out on possible imprecise mental maps of local and/or global commanders).

Insofar as complex adaptive systems can be regarded as being essentially open-ended problem-solvers, their lifeblood consists mostly of novelty. The ability of a complex adaptive system to survive and evolve in a constantly changing environment is determined by its ability to continually find – either by chance, or experience, or more typically both – insightful new strategies to increase its overall "fitness" (which is, of course, a constantly changing function in time).

Military campaigns likewise depend on the creative leadership of their commanders, success or failure often hinging either on the brilliant tactic conceived in the heat of combat or the mediocre one issued in its place.

To be realistic, such novelty must not consist solely of a randomly selected option from a main-options list – which is a common approach taken by conventional warfare models – but must at least have the possibility of being as genuinely unanticipated in the model as it often is on a real battlefield. To this end, each command-agent (and to a somewhat more limited extent, each ISAACA) must possess both a memory and an internal anticipatory mechanism which it uses to select the optimal tactic and/or strategy from among a set of predicted outcomes. This is an important point: except for doctrine and the historical lessons of warfare, the super-set of tactics *must not be hard-wired in*.


Such local rule-based agent-simulations are well suited for

- studying the general efficacy of combat doctrine and tactics
- exploring emergent properties and/or other "novel" behaviors arising from low-level rules (even doctrine if it is well encoded²⁸)

²⁸ It is an intriguing speculation that doctrine as a whole may contain both desirable and undesirable latent patterns that emerge only when allowed

- capturing universal patterns of combat behavior by focusing on a reduced set of critical drivers
- suggesting likelihood of possible outcomes as a function of initial conditions
- use as training tools along the lines of some commercially available agent-based "games" such as SimCity, SimFarm and SimLife²⁹
- providing near-real-time tactical decision aids by providing a "natural selection" (via *genetic algorithms*; see Part I [28], page 93) of superior tactics and/or strategies for a given combat situation
- giving an intuitive "feel" for how and/or why unanticipated events occur on the battlefield, and to what extent the overall process is shaped by such events

Figure 19. Schematic representation of a ISAACA simulation



Ideally, one would hope to find universal patterns of behavior and/or tactics and strategies that are independent of the details of the makeup of individual ISAACAs.

to "flow" through a system of elementary agents. An agent-based model of combat may provide an ideal simulation environment in which to explore such possibilities.

²⁹ W. Wright, *SimCity* (computer game), Orinda, California: Maxis Corporation, 1989.

Agent-based simulations ought *not* be used either to predict real battlefield outcomes or to provide a realistic simulation of combat. While commercial networkable 3D virtual-reality games such as DOOM³⁰ are much better suited to providing a virtual combat environment for training purposes, agent-based simulations are designed to help understand the basic processes that take place on the battlefield. It is not realism, for its own sake, that agent-based simulations are after, but rather a realistic understanding of the drivers (read: interactivity, decision-making capability, adaptability, and so on) behind what is really happening.

*Swarm*³¹

Swarm is a multi-agent simulation platform for the study of complex adaptive systems. It is currently under development at the Santa Fe Institute.³²

The goal of the *Swarm* project is to provide the complex systems theory research community with a fully general-purpose artificial-life simulator. The system comes with a variety of generic artificial worlds populated with generic agents, a large library of design and analysis tools and a "kernel" to drive the actual simulation. These artificial worlds can vary widely, from simple 2D worlds in which elementary agents move back and forth to complex multi-dimensional "graphs" representing multidimensional telecommunication networks in which agents can trade messages and commodities, to models of real-world ecologies.

Swarm has been intentionally designed to include as few ad-hoc assumptions about the design of a complex system as possible, so as to provide a convenient, reliable and standardized set of software tools that can be tailored by researchers to specific systems.

Though the prototype has been written using the C programming language, it is object-oriented in style. Future

³⁰ *Id Software*, World-Wide Web URL link = <http://www.idsoftware.com>.

³¹ This section is based on the papers "An Overview of the *Swarm* simulation system," by '94 *Swarm* Team, Santa Fe Institute and "The *SWARM* simulation system and individual-based modeling," by D. Hiebler.

³² Working papers and other documentation about the project can be found at the *SWARM* site at the Santa Fe Institute: World-Wide-Web URL link = <http://www.santafe.edu/projects/swarm/>.

versions of Swarm will be implemented using the Objective-C language. Objective-C is an object-oriented extension of the C language that is widely available as part of the GNU C compiler, and is available on the World-Wide-Web.

Everything in Swarm is an object with three main characteristics: *Name*, *Data* and *Rules*. An object's *Name* consists of an ID that is used to send messages to the object, a type and a module name. An object's *Data* consists of whatever local data (i.e. internal state variables) the user wants an agent to possess. The *Rules* are functions to handle any messages that are sent to the object. The basic unit of Swarm is a "swarm": a collection of objects with a schedule of event over those objects. Swarm also supplies the user with an interface and analysis tools.

The most important objects in Swarm, from the standpoint of the user, are *agents*, which are objects that are written by the user. Agents represent the individual entities making up the model; they may be ants, plants, stock brokers, or combatants on a battlefield. *Actions* consist of a message to send, an agent or a collection of agents to send the message, and a time to send that message. Upon receiving a message, agents are free to do whatever they wish in response to the message. A typical response will consist of the execution of whatever code the user has written to capture the low-level behavior of the system he is interested in. Agents can also insert other actions into the schedule.

Three other properties of Swarm are noteworthy:

1. *Hierarchy*. In order to be better able to simulate the hierarchical nature of many real-world complex systems, in which agent behavior can itself be best described as being the result of the collective behavior of some swarm of constituent agents, Swarm is designed so that agents can be themselves be swarms of other agents. Moreover, Swarm is designed around a time hierarchy. Thus Swarm is both a nested hierarchy of swarms and a nested hierarchy of schedules.
2. *Parallelism*. Swarm has been designed to run efficiently on parallel machine architectures. While messages within one swarm schedule execute sequentially, different swarms can execute their schedules in parallel.

3. *Internal Agent Models.* One can argue that agents in a real complex adaptive system (such as the economy) behave and adapt according to some internal model they have constructed for themselves of what they believe their environment is really like. Sometimes, if the environment is simple, such models are fixed and simple; sometimes, if the environment is complex, agents need to actively construct hypothetical models and testing them against a wide variety of assumptions about initial states and rules and so forth. Swarm allows the user to use nested swarms to allow agents to essentially create and manage entire swarm structures which are themselves simulations of the world in which the agents live. Thus agents can base their behavior on their simulated picture of the world.

Among the many kinds of problems that Swarm is well suited for are economic models (with economic agents interacting with each other through a market), the dynamics of social insects, traffic simulation, ecological modeling, simulation games such as SimCity and SimLife, and general studies of complex systems, cellular automata, and artificial life.

Judging from its specifications, Swarm is ideally suited to act as the backbone of a full system-level simulation of land combat. However, because of the unproven status of Swarm itself (it is, of this writing, in the final stages of beta-testing), such a project requires a considerable investment of both risk-tolerance and commitment.

Tier VIII: Original Conceptualizations of Combat

Tier VIII represents the potentially most exciting – and certainly most far-reaching – tier of the eight tiers of applications. It consists of using complex systems theory inspired ideas and basic research to develop *fundamentally new conceptualizations of combat*. It asks what is simultaneously the most direct and most expansive possible question regarding the complex systems theoretic view of land combat: "What are the universal characteristics of land combat, thought of as a complex adaptive system?" Because of the very speculative nature of this question, *Tier-VIII* thus also necessarily takes the longest-term view of expected development time. But while this tier obviously entails the greatest risk, it also promises to yield the greatest potential payoff.

Tier-VIII complements *Tier-IV*, on which the objective is to find complex systems theoretic "measures" that describe combat. *Tier-VIII* is concerned with what to do with those measures once they are found.

Ideally, complex systems theory may suggest ways in which battlefields must be configured (or compelled to self-organize) to be maximally adaptable to the most wide-ranging set of environmental circumstances. It would be most interesting, for example, to be able to determine what doctrine, constraints and/or specific rule sets prescribing what local actions can and cannot be taken that are most conducive to pushing a combat force closer to the edge-of-chaos?

In the remainder of this section we discuss briefly a few speculative ideas and technologies that might be used to develop applications lying on this tier of applicability. The discussion is mostly qualitative and is designed to plant seeds for future work. Some more speculative and open questions, whose answers undoubtedly require work to be done on this tier, appear in the concluding section of this paper.

Dueling Parasites

Genetic algorithms have thus far figured very prominently on a variety of tiers of applications, ranging from helping design more efficient and robust command and control structures on *Tier-II* to acting as the source of the "adaptive intelligence" of adaptive autonomous agents in a multi-agent simulation of combat on

Tier-VII. The reason for this, of course, is that genetic algorithms are a mainstay of most complex systems theory models.³³ Here we outline a potentially powerful generalization of the basic genetic algorithm introduced by Hillis [26], which may have a natural application to the modeling of combat.

Conventional genetic algorithms search for "solutions" to problems by "evolving" large populations of approximate solutions, each candidate solution represented by a chromosome (see Part I [28], pages 93-101). The genetic algorithm evolves one population of chromosomes into another according to their fitness using various genetic operators (such as crossover and mutation), and, eventually, after many generations, the population comes to consist only of the "most-fit" chromosomes.

This basic recipe has, of course, been shown to be useful for finding near-optimal solutions for many kinds of problems. One of the major difficulties that *all solution schemes* for solving combinatorial optimization problems must contend with, however, is the classical problem of the search space containing local optima: once a search algorithm finds what it "thinks" is the global optimal solution, it is generally difficult for it to find ways to not be "locked into" the local optimum.

Hillis attacks this problem by exploiting host-parasite interactions among two coupled genetic algorithm populations. To illustrate the idea, consider his testbed system, which consists of finding a sorting algorithm for elements of a set of fixed size that requires the smallest number of comparisons and exchanges to be made among the elements. The overall problem is to design an efficient *sorting network*, which is a sorting algorithm in which the sequence of comparisons and exchanges is made in a predetermined order. A candidate sorting network, once defined (by a chromosome), is easy to test.

Now, Hillis' idea is to set up not one but *two* interacting genetic algorithm populations, one population consisting of "solutions," or sorting programs (the *hosts*), and the other consisting of "sorting problems" (the *parasites*). Having the two populations interact effectively sets up an "arms-race" between the two populations. While the hosts are trying to find better and better ways to sort the problems, the parasites are trying to make the

³³ A cautionary note about a too-cavalier use of genetic algorithms is sounded in Appendix E.

hosts less and less adept at sorting the problems by making the problems "harder."

The interaction between the two populations dynamically alters the form of the fitness function. Just as the hosts reach the top of a fitness "hill," the parasites deform the fitness landscape so that the hill becomes a "valley" that the hosts are then forced to find ways to climb out of and start looking for new peaks. When the population of programs finally reaches a hill that the parasites cannot find a way to turn into a valley, the combined efforts of the co-evolving hosts and parasites has found a global optimum. Thus, the joint, coupled, population pools are able to find better solutions quicker than the evolutionary dynamics of populations consisting of sorting programs alone.

The application to combat modeling is conceptually straightforward. The idea is to apply genetic algorithms not to just one side of a conflict, or to use genetic algorithms to find "optimal" combat tactics for fixed sets of constraints and environments, but to use *joint, coupled, pools of populations*, one side of which represents a set of tactics or strategies to deal with specific scenarios, and the other side of which seeks ways to alter the environment in ways that make it harder and harder for those tactics or strategies to work.

Percolation Theory and Command and Control Processes

Percolation theory represents the simplest model of a disordered system. Consider a square lattice, where each site is occupied randomly with probability p or empty with probability $1-p$. Occupied and empty sites may stand for very different physical properties. For simplicity, let us assume that the occupied sites are electrical conductors, the empty sites represent insulators, and that electrical current can flow between nearest neighbor conductor sites. At low concentration p , the conductor sites are either isolated or form small clusters of nearest neighbor sites. Two conductor sites belong to the same cluster if they are connected by a path of nearest neighbor conductor sites, and a current can flow between them. At low p values, the mixture is an insulator, since a conducting path connecting opposite edges of the lattice does not exist. At large p values, on the other hand, many conduction paths between opposite edges exist, where electrical current can flow, and the mixture is a conductor. At some concentration in between, therefore, a threshold concentration p_c must exist where for the first time electrical

current can percolate from one edge to the other. Below p_c , we have an insulator; above p_c we have a conductor. The threshold concentration is called the percolation threshold, or, since it separates two different phases, the critical concentration. The value of the critical concentration depends on the connectivity pattern of the lattice.

What does this have to do with command and control structures and processes? Conceptually, one can form analogies between conductor sites and information processing and/or data-fusion centers and between electrical current and information. Information can be radar contact reports, commands pumped down echelon or raw intelligence data. The problem of determining the efficacy of a given flow of information can be solved by interpreting it as a percolation problem. Among the intriguing questions that inevitably arise from drawing such an analogy, are "What command and control architectures are most conducive to information flow?", and "What are the inherent vulnerabilities in the existing command and control structure?", and so on.

Woodcock and Dockery also liken percolation through a lattice to the percolation of military forces through an area of obstacles or a combat zone of deployed adversarial forces.³⁴

There is a lot of interesting theoretical work that is being done on *random graph theory* [46], which deals with how the global topological properties of a mathematical graph³⁵ (such as its overall connectivity, the maximum number of connected clusters of sites that it contains, and so on) change as a function of the number of nodes and links in the graph. Theoretical results such as these provide important information about how the overall efficacy of, say, a command and control communications network depends on quantifiable measures of its topology.

Exploiting Chaos

Deterministic chaos seems grounded in paradox: "simple" equations generate "complicated" behavior, "random" appearing trajectories harbor embedded "patterns," and so on. Many potential applications depend heavily on apparent paradoxes (if

³⁴ See reference [18], page 323.

³⁵ A mathematical graph can be thought of as the network describing an Integrated Air Defense System, but where the presence or absence of a given link between nodes is specified by a probability distribution.

not outright oxymoronic assertions) such as these, seeking to find ways to exploit the inherent regularities that systems exhibiting a deterministic form of chaos are naturally predisposed to possess. Here we mention three such seeming paradoxical properties of chaotic systems that might be exploited on both practical and theoretical levels:

- *Chaotic Control*
- *Chaotic Synchronization*
- *Taming Chaos*

Chaotic Control

Chaotic control has been discussed in Part I [28].³⁶ It refers to using a chaotic system's sensitivity to initial conditions to stabilize regular dynamic behaviors and to effectively "direct" chaotic trajectories to desired states. It has been amply demonstrated both theoretically and practically for a wide variety of real physical systems. It is interesting to note that this is a capability that has no counterpart in nonchaotic systems for the ironic reason that the trajectories in nonchaotic systems are stable and thus relatively impervious to desired control.

From a theoretical point of view, chaotic control could conceivably be used by decision makers to *selectively guide, or "nudge," combat into more desired states*. Of course, this presupposes that an appropriate phase-space description of combat has been developed, and all of the relevant control parameters have been identified.

Chaotic Synchronization

Like chaotic control, the idea of being able to synchronize coupled chaotic systems seems almost an oxymoron, but has its roots in the same basic idea of selectively driving a chaotic dynamical system to restrict its motion to a desired subspace of the total phase space. Chaotic synchronization refers to selectively coupling two identical chaotic systems in such a way that they then evolve with their corresponding dynamical variables exhibiting exactly the same behavior in time

³⁶

See pages 59-63.

First introduced by Pecora and Carroll [47], the underlying principle is to look for a range of parameter settings for which the joint phase space of two chaotic systems is stable to motion on a subspace where the motion is either regular or is of another type of chaotic behavior [1].

Theoretical analysis of the general question of what happens when one chaotic dynamical system is used to "drive" another – which is how, on a conceptual level, one can interpret the selective "nudging" of forces on a battlefield – has potentially enormous implications for our ability to predict how the overall system of ground forces will react. Synchronization also has clear applications to communications and developing a robust and reliable form of IFF. Moreover, the fact that a generalized relationship between driving signals and response system signals *exists at all*, suggests that this function can in principle be found and used for prediction purposes.³⁷ Careful attention to the theory behind, and potential practical applications of, synchronized chaos is likely to have a high payoff.

Taming Chaos

A very recent addition to the list of counterintuitive behaviors of chaotic systems is what can be described as "taming chaos" with chaos. Disorder and noise in physical systems typically tends to destroy any existing spatial or temporal regularities, or so one's intuition would lead one to expect. Not so! For example, it can be shown that some nonlinear systems are able to transfer information more reliably when noise is present than when operating in a noiseless environment.³⁸

Braiman, Lindner and Ditto³⁹ have also recently reported an interesting experiment in which an array of periodically forced pendula lapses into spatiotemporal chaos when the pendula are identical, but then snaps into a periodic behavior when disorder is *added* to the system! Braiman, *et. al.*, speculate that disorder can be used to tame spatiotemporal chaos and suggest that "the role of disorder in spatially extended systems may be less of a randomizing influence than an intrinsic mechanism of pattern formation, self-organization and control."⁴⁰ Again, the ability to selectively alter the apparently chaotic patterns of behavior on

³⁷ See reference [1], pages 163-166.

³⁸ This is a phenomenon called *stochastic resonance*, see, for example, F. Moss and K. Weisenfeld, *Nature*, Volume 373, 1995, 33-36.

³⁹ Y. Braiman, J. F. Lindner and W. L. Ditto, *Nature*, 30 Nov 1995, 465.

⁴⁰ *ibid.*, page 467.

the spatio-temporal arena of the battlefield has broad implications and potentially an enormously high payoff. Whether, or to what extent, such a "taming" of chaos is possible requires us to first carefully study the general phenomenon for simple models of combat.

Pattern Recognition

"If you see a whole thing - it seems that it's always beautiful. Planets, lives....
But close up a world's all dirt and rocks. And day to day, life's a hard job, you
get tired, you lose the pattern." - Ursula K. LeGuin

What is "Battlefield Intuition"? In a previous section (see page 26) we suggested that battlefield intuition is an innate ability to perceive (though perhaps not to articulate) underlying patterns in what otherwise seems to be irregular behavior. We compared it to the intuition of the successful stock-broker on wall-street, who has an intuitive "feel" for when certain stocks will rise and fall. Whatever the underlying basis is for battlefield intuition, however, certainly one of the most important fundamental problems facing any commander is the pattern recognition problem. In order to make sound decisions a commander must know what is *really* happening on the battlefield. "Knowing what is really happening" does not just mean finding better ways to get at ground truth; it means seeing *patterns of behavior* that others have either not looked for or have simply missed seeing altogether. This is also a fundamental problem faced by any agent in a complex systems theoretic multi-agent based simulation of, say, a natural ecology. In order for an agent to survive and successfully evolve in the ecology, it needs to identify the parts of the environment that are relevant and understand how the relevant parts *really* fit together, not how they appear to fit. While solving the pattern recognition may, at first, appear to have little in common with "complex systems theory," it is in fact a problem that lies at the core of any complex systems theoretic approach. It is thus also lies at the core of a complex systems theory approach to land warfare.

Now, while this is not to say that complex systems theory has "solved" the pattern recognition problem, it is meant to suggest that some of the tools that complex systems theory has developed for dealing with the general pattern recognition problem *can* also be applied to discerning patterns on the battlefield.

The "conventional" tool-kit for dealing with patterns embedded in otherwise chaotic dynamics comes from nonlinear dynamics and consists of four basic parts [1]:

1. *Finding the signal*, in which the signal of interest is first extracted from the raw data; of course, in many instances, the raw data may be the signal, since there is no a-priori way of discerning noise from meaningful information
2. *Finding the phase space*, which consists of the time-delayed embedding technique of creating a d-dimensional vector out of an a-priori "list" of numbers
3. *Classifying the signal*, which can be done by using such measures as Lyapunov exponents, various fractal dimensions, and other quantities independent of the initial conditions
4. *Developing a model and predicting future behavior*, based on the classifications made during the previous step

We will not go into any greater detail about any of these steps, except to say that these are techniques that have by now been fairly well established in the research literature. Of course, as with any general set of tools, each of the tools in this tool-chest has certain advantages and disadvantages and is more or less adept at dealing with specific kinds of data.

In addition to these more or less "conventional" tools borrowed directly from nonlinear dynamics theory, however, there are other – more theoretical and speculative – methods available. We mention three such methods: (1) high-level rule extraction using genetic algorithms, (2) self-organizing neural nets to sort raw information, and (3) data-base mining for knowledge.

High-Level Rule Extraction

Richards, Meyer and Packard [50] have recently suggested a way to extract two-dimensional cellular automaton rules directly from experimental data. Recall that two-dimensional cellular automata are a class of spatially and temporally discrete, deterministic dynamical systems that evolve according to a local evolutionary rule.⁴¹ Richards, *et. al.*'s idea is to use a genetic algorithm to

⁴¹ For a review of cellular automata, see pages 81-91 in Part I [28].

search through a space of a certain class of cellular automata rules for a local rule that best reproduces the observed behavior of the data. Their learning algorithm (which was applied specifically to sequential patterns of dendrites formed by NH_4Br as it solidifies from a supersaturated solution) starts with *no a-priori* knowledge about the physical system. It, instead, builds increasingly sophisticated "models" that reproduce the observed behavior.

Though Richards, *et. al.*'s NH_4Br testbed has *a-priori* little to do with combat, it is in principle not that far away. Like combat, dendritic NH_4Br data exhibits pattern structure on many different length scales and a dynamics takes place on different time scales. Moreover, there is often very little information available regarding the physical variables describing the dendritic solidification of NH_4Br . While one can determine whether a given point is solid or liquid, for example, one typically knows nothing about the solute concentration or temperature field in the liquid. The situation is much the same in combat, where one may know the disposition of one's forces and perhaps something about what individual combatants are doing at what time, but the specifics of their actions and of any internal dynamics they are following are effectively unknown. In the NH_4Br , despite this lack of knowledge of what is happening on the micro-level, Richards, *et. al.*'s algorithm is able to find a rule that qualitatively reproduces the observed data.

Richards, *et. al.* comment that while the exact relationship between the rule found by their genetic algorithm and the fundamental equations of motion for the solidification remains unknown, it may still be possible to connect certain features of the learned rule to phenomenological models.

"We propose that this type of 'derivability gap' is the rule, rather than the exception for most complex spatial patterns observed in nature. For such phenomena, it may be impossible to derive models which explain observed spatiotemporal complexities directly from fundamental equations and 'first principles.' Though perhaps underivable, the dynamical structure extracted by the learning algorithm is undeniable, and represents a new type of progress, perhaps the primary kind of understanding possible for complex patterns."⁴²

It is tempting to speculate what insights a similar approach to extracting "low-level rules" from "high-level observed behavior" on the battlefield might have to offer.

Self-Organizing Maps

A Self-Organizing Map (SOM) is a general unsupervised neural-network.⁴³ Introduced by Kohonen⁴⁴ in the early 1980s, it is designed to order high-dimensional statistical data so that inputs that are "alike" generally get mapped to each other. Unlike backpropagating neural nets, that require that the output part of a desired input-output set of pairs is known a-priori, unsupervised learning effectively *tells the trainer* what latent patterns and similarities exist within a block of data. Thus, it can be used as a means by which to "self-organize" ostensibly patternless masses of information like raw intelligence data, or existing databases consisting of various unstructured bits of information about an adversary. The idea is literally to allow the raw data to "tell the intelligence analyst" what kinds of innate structural patterns might exist in the data. It is not a cure-all – as it requires behind-the-scenes preprocessing and some assumptions to be made about what kind of structuring and "document-distance" measures are appropriate – but the methodology potentially provides an important first step in helping an analyst, or field commander, intelligently sift through apparently reams of formless information.

An example of how SOMs can be used as "information organizers" is a recent effort called WEBSOM.⁴⁵ WEBSOM is designed to automatically order, or organize, arbitrary free-form textual information into meaningful maps for exploration and search. It automatically organizes documents into a two-dimensional grid so that the closer two documents are "related" to each other the closer they appear together on the grid. More specifically, WEBSOM has been applied to ordering documents on the World-Wide-Web (WWW).

Anyone who has spent even a short time "cruising" the WWW knows that while there is a tremendous amount of information available on the web, *desired* information is more often than not extremely difficult to find. Web search engines such as *Lycos*⁴⁶ or *AltaVista*⁴⁷ are "intelligent" enough to retrieve some meaningful sites for specific queries, but are next-to-useless when it comes to finding sites or files in cases where the actual subject or object of

⁴³ See Part I [28], pages 116-130.

⁴⁴ T. Kohonen, *Proceedings IEEE*, Volume 78, 1990, 1464-1480.

⁴⁵ The best available information on WEBSOM can be found at the WWW address <http://websom.hut.fi/websom/>.

⁴⁶ <http://lycos.cs.cmu.edu/>.

⁴⁷ <http://altavista.digital.com>.

interest is only vaguely known. Even in those cases where existing search engines are able to find a few useful files, these files are often buried deep in an otherwise lengthy list of files that are only marginally related to a specific query, if at all.

WEBSOM is designed to help such free-form searches by automatically organizing a set of documents so that related documents appear close to each other. An initial testbed for the technique consisted of 4600 full-text documents from the "comp.ai.neural-nets" newsgroup, containing a total 1,200,000 words.⁴⁸

After being organized by WEBSOM, the newsgroup documents can be viewed on four levels. The top level consists provides an overview of whole document collection. It consists of individual nodes representing the highest-level clusters of documents, arranged by similarity, and uses grey-scales to indicate clustering density. Levels two and three are accessed by clicking the mouse on a desired super-cluster of related documents, and represent successively deeper nestings of documents organized into mid-level clusters. The fourth, and final, level is accessed by clicking the mouse on a desired cluster on the third level, and consists of actual document listings, now grouped such that all "nearby" documents are closely related. As one proceeds down from the top-most to bottom-most level, one goes from the most general clusters (neural nets, fuzzy logic, forecasting,...) down to more finely divided clusters (neural nets in plant manufacturing, fuzzy control of neural nets, ...) down to individual documents. WEBSOM thus effectively maps out the entire "document space" according to what documents actually inhabit that space. "Closeness" is interpreted with respect to semantic content, as approximated by a statistical sampling of word contexts. Other measures could be devised for other applications.

In this specific newsgroup example, WEBSOM provides a display of the similarity relations of the subject matters of the documents. These are reflected in the distances between documents in the document map. The density of documents in different parts of the map are reflected by varying shades of grey on the document display. One can easily imagine suitably generalized versions of this methodology being applied to organizing raw intelligence data. More speculatively, one can

⁴⁸ An interactive demonstration of using WEBSOM for this example appears at the site <http://websom.hut.fi/websom/>. The discussion follows the paper "Newsgroup exploration with WEBSOM method and browsing interface," by T. Honkela, *et. al.* that can be retrieved from this site.

imagine using SOMs to provide "unconventional" partitionings of battlefield processes. That is, just as WEBSOM is able to tell us something about the natural ordering in document space, from which we are then able to infer patterns that can be used to more intelligently guide our search for information, so SOMs may be able to tell an analyst or field commander something about the natural ordering in "combat space," from which an analyst or field commander is then able to infer patterns that he can use to make "more informed" decisions.

Data-Base Mining for Knowledge

Frawley, et. al. [] define *knowledge discovery* as the "nontrivial extraction of implicit, previously unknown, and potentially useful information from data." Much work has recently been done in the area of database mining, which is essentially an application of the scientific method to database exploration. The basic problem is easy to state: given a data set find a pattern, or patterns, that describes meaningful, consistent relationships among subsets of the database. Ideally, of course, the pattern should be simpler to articulate than merely enumerating all the facts. Knowledge discovery is therefore generally concerned with inducting, from data, possible rules or "laws" that may have been responsible for generating that data. The connection to the basic pattern recognition problem on the battlefield, should again be obvious from a complex systems theory point of view: given that a database D contains a "record" of a land warfare campaign, we are interested in finding the "implicit, previously unknown, and potentially useful information" that can be extracted from D.

We do not have the space here to go into the details of the many techniques that are available for addressing the general database mining problem. We briefly mention three recent examples []:

- *Kepler*, which is a system designed to find functional relationships among quantitative data. Applied to a data base consisting of experimentally derived fluid flow data, for example, Kepler is able to "discover" such basic laws as Bernoulli's theorem for laminar flow.
- *Thought*, which is capable of incrementally discovering production rules by classifying and abstracting from given examples, and then finding implications between the descriptions according to the relationships it finds among corresponding clusters of data.

- *Posch*, which is an automated artificial intelligence systems designed to discover causal relationships in a medical-record database

Some, though not all, data-mining tools can be thought of as being natural-language equivalents of the more number-intensive techniques developed by nonlinear dynamics for finding underlying patterns in number fields. It would be an interesting exercise to use some of the available data-mining techniques to explore what hidden relationships might exist in historical combat data, for example, not to mention using such techniques for exploring patterns and relationships in data that summarizes combat exercises and /or actual combat.

Fire-Ant Warfare

As an example of a potentially far-reaching technology that is clearly inspired by complex system theoretic concepts is the *Fire-Ant Warfare* idea recently put forth by Libicki [38]:

"Today, platforms rule the battlefield. In time, however, the large, the complex, and the few will have to yield to the small and the many. Systems composed of millions of sensors, emitters, microbots and miniprojectiles, will, in concert, be able to detect, track, target, and land a weapon on any military object large enough to carry a human. The advantage of the small and the many will not occur overnight everywhere; tipping points will occur at different times in various arenas. They will be visible only in retrospect."

The idea is to exploit the collective intelligence of a swarm of (perhaps thousands of) tiny intelligence-gathering machines and small smart-weapons. Libicki suggests that systems of millions of sensors, emitters, microbots and miniprojectiles can be used in concert to detect, track, target and land a weapon on military targets. This approach is reminiscent of Rodney Brooks' [11] micro-bot artificial-life approach to artificial intelligence. In contrast to the traditional top-down methods that emphasize abstract symbol manipulations and high-level reasoning skills, Brooks proceeds from the bottom-up by using many small and individually "simple" micro-bots, or autonomous agents, to assemble a "collective intelligence" that – when the agents act in concert – is capable of performing very sophisticated tasks. We will not say more about this very active school of research, except to suggest that artificial-life-like fire-ant warfare represents not just a conceptual advance in the way we think about warfare, but a significant technological advance in how we conduct it as well.

Suggested Directions

The most important overall suggestion that can be made regarding the applicability of complex systems theory to land warfare is to *be patient!* As discussed at length in Part I of this report, and stressed repeatedly throughout both volumes, complex systems theory is a very young, very immature science, which – at this time – is not even sure of its *own* future direction, much less of its applicability to other, specific areas. Ironically, we must therefore find ways to "tune" ourselves to the "edge-of-chaos" (see Part I, page 76), and be on special guard against both seeing in complex systems theory more than is currently there, and against prematurely walking away in disappointment from what we have not taken the proper time to discover is there already.

General Directions

General guidelines for applying some of the basic lessons learned from nonlinear dynamics and complex system theory include:

- **Familiarization at all Levels of the Military Chain-of-Command.** If the ideas of nonlinear dynamics and complex systems theory are to percolate through all levels of the military, first and foremost it is important for its leaders to come up to speed on, and develop an intuition for, some of the technical aspects of these approaches. More simply, *one has to get one's hands dirty!* There is no better way of getting a feel for why chaos and complexity have potentially so much to tell us about combat, than by sitting down in front of a computer and "playing" with a few simple but well-chosen models. The basics of nonlinear dynamics and complex systems theory must also be taught at an early stage at military schools.
- **Develop "Nonlinear Intuition."** It is vital for every decision maker to go beyond the conventional "linear" intuition and develop an intuition for the kinds of nonlinear behaviors pervasive in complex systems. Arguably, our most successful battlefield commanders already possess it. But, like the mysteriously successful stock-brokers who have an uncanny "feel" for which way the stock-market will turn but are unable to explain exactly where their feeling comes from, one strongly suspects that only a very few field commanders

possessed of an uncanny "battlefield intuition" can exactly articulate the source of their intuition. Now, while "battlefield intuition" is certainly not something that can easily be taught (or taught at all – there is no mold for commanders like Patton!), there is a part of it that just as certainly rests on an appreciation of, and intuition for, how patterns can form in nonlinear dynamical systems; and this is something that, with proper instruction and practice, *can* be learned.

- **Emphasize Strong Interdisciplinarity.** If there is one universally agreed upon "insight" that has emerged out of *Santa Fe Institute's* first dozen years of existence it is that progress in CST demands an interdisciplinary approach. It is not enough to merely know everything there is to know about some complex system. One must also be either well versed in many other disciplines, or, ideally, be in the company of and continually engage in a free-flow of ideas with individuals who are well versed in many other disciplines. Complex adaptive systems, it seems, are best studied by other complex adaptive systems. A cross-fertilization of ideas and approaches is absolutely vital for the largely inductive process by which progress is made in CST, and can come about only in an open interdisciplinary setting.
- **Look for Inherent Nonlinearities in Conventional Models.** A fundamental lesson of nonlinear dynamics theory is that one can almost always expect to find some manifestation of chaos whenever nonlinearities are present in the underlying dynamics of a model. This fundamental lesson has potentially significant implications for even the simplest combat models. Though some work has recently been done to determine the implications of having nonlinearities embedded within conventional models, many important insights into how our current models of land combat really behave remain to be discovered.
- **Redefine Traditional MOEs and Data Collection Requirements.** If land combat is a bona fide candidate system for study as a complex system it must, initially, be treated essentially as "just another system" for study by CST. This means that the first real research task is to re-examine what we know about the conduct of war from the perspective of complex systems theory. The traditional

military theorist's point of view, with a predisposition towards force-strength, fire-power, attrition statistics, and so on, may be inappropriate, incomplete and/or simply inadequate to describe combat from a CST point of view. CST will likely redefine basic data collection requirements and suggest that we rethink our answers to many fundamental questions: *What measures are appropriate? What data are missing?* Satellites, for example, which have not heretofore been thought of combat "data collectors," can be gainfully employed to obtain measures of the "overall flow" of combat.

- **Start with "Minimal Idea Models," not Full-Blown Detailed System Models.** Conventional military wisdom expects the justification for an important decision to come from a sufficiently "fancy model." The more lines of computer code a model has, so the conventional wisdom goes, the more attention should be given to what the model has to say. As an important counterpoint, it is well worth remembering that some of the greatest breakthroughs in physics have come from very simple models that capture only the essence of a system. The hydrogen atom, for example, which is the simplest atom, was the key "model" that led to the development of quantum mechanics. The first task of any fundamental research effort – and this is what finding ways of applying complex systems theory to land combat must necessarily be viewed as – is to find a simple enough system that, while it is not an exact replica of the system that one is trying to understand and may lack many of its real-world complications, is able to capture some of the essential properties of the real system.
- **Attack Problem from Diverse Fronts.** Complex system theory consists of any and all ways of going about understanding the behavior of a complex system. This means attacking the problem from diverse fronts. The previous bullet suggests that the first goal of *any* fundamental research effort is to develop a simple model that captures the essential behavior of a system. A second, and equally as important, goal is to try to understand lots and lots of different systems in the hopes of discovering some universal patterns of behavior that are common to all of them. This second goal is really how "complex systems theory" is practiced at the *Santa Fe Institute*. Given that other "simple models" of other complex systems are being

developed by researchers in other disciplines, it is incumbent on the military research and development community to contribute its own "simple models" of combat to this growing set of exemplar models of complex systems. The benefit would be two-fold: (1) behaviors of models of combat could be directly compared with models of other complex systems to help discern any universal patterns, and (2) the simple model can be studied on its own terms from a complex systems theory perspective to lend insight into the essential behaviors on a battlefield. A third goal, or strategy, is purely empirical. We need to develop new tools to record relevant data (and to re-examine historical data) from a complex systems theory perspective.

Some Open Questions

We conclude this paper by presenting a list of some open and overtly speculative questions:


Are there measures of combat "complexity"?

Are there quantifiable measures of complexity on the battlefield that can be used as indicators of qualitative changes in such properties as force strength, battle tempo, morale, etc.? Are there measures with which the overall "fighting health" can be ascertained and/or predicted?

Can patterns of observed chaotic data be exploited?

Suppose that one can show that certain aspects of land combat are "chaotic" (in the technical sense). The basic question to ask is whether we can exploit any underlying order in the deterministic chaos to make short-term predictions? More precisely, we ask, "Given data set X observed during time period ΔT , how do you predict what the system will be doing ($=\phi(X)$) during time Δt at a time τ later?" (see figure 20). Can we use historical attrition data? What other kinds of data might prove useful? How long can τ be before $\phi(X)$ loses its predictive value? What are the parameters describing the behavior of land combat that are most amenable to short-term predictability

Figure 20. Schematic of basic chaotic-pattern exploitation question



What is an appropriate phase space description for combat?

What parameters can be used such that, in their phase space, a strange-attractor-like pattern emerges on a higher level? Suppose that a general combat "configuration" – including force structures, tactics and strategies – were encoded as a genome. The genome would thus act as a reference template for all possible actions, counteractions, counter-counteractions, and so on, of all possible combat configurations. Genetic algorithms could be used to search for possible parameter spaces for which the long-term behavior of the system, as a whole, appears more "ordered."

Can the chaos of combat be "controlled" or "tamed"?

Part I [28] of this paper discussed a relatively new technique called *chaotic control*.⁴⁹ Chaotic control is a technique whereby the extreme sensitivity of chaotic systems to small perturbations to initial conditions (the so called "butterfly effect") is exploited to stabilize regular dynamic behaviors and to effectively "direct" chaotic trajectories to a desired state. The question is, "Are there ways to 'nudge' the seemingly erratic patterns (i.e. trajectories) of behavior on the battlefield to follow desired courses?" If an appropriate phase space description of combat can be found (see preceding question), it might be possible to map out the behavioral characteristics of various regions of that phase space. For example, can we find ways to selectively apply feedback to induce desired behavioral transitions? In a combat setting,

⁴⁹ See Reference [1], pages 59-62.

control parameters might include the character, makeup and numerical force strength, type of weaponry, timing and accuracy of intelligence reports, tactics, etc.

What are the "optimal" strategies of adaptation on the battlefield?

Can complex systems theory be used to suggest ways in which a combat system can continue to best adapt to changing situations and scenarios? Given a system $S(E)$ that has successfully adapted to environment E , suppose E changes to E' . What must S do to itself to get to a state S' (in which it has adapted to E') and how fast must it do it?

What role does the psychology of the individual combatant play in shaping the combat process?

One of the first things a complex systems theorist does while developing a model of, say, a natural ecology, is to understand the dynamics between individual organisms and their environment. Yet, from a military modelist's point of view, the psychology of the individual combatant, and the role the combat environment plays in shaping his behavior, remains the one area that is almost never developed to any depth. This is not to suggest that the problem is easy; it is not. It is meant to suggest that any earnest attempt at describing the combat environment from a complex systems theory point of view must carefully consider the relationship between combatant and environment.

How "complex" must a combat system be in order for it to be amenable to the tools of complex systems theory?

This applies to both spatial and temporal scales. Two individual combatants, though they are both complex organisms with rich inner psychological dynamics, do not make up a "complex system." Are ten combatants enough? Are a hundred?

How can one quantify the true value and nature of "information" on a battlefield?

A popular current buzzword in the military is "information warfare," which deals with the exploitation of various forms of information – textual, electromagnetic, psychological, and so forth. This question deals with a more fundamental meaning of the word "information." Whereas the traditional interpretation of

combat has been strictly materialist – combat as sets of material objects reacting to other material objects – complex systems theory views combat as a system consisting of material objects reacting to information – information about the perceived state of the system, information about the prior expectations of an opponent's force structure and ability to fight, and so on. Combat is understood not just as a set of local fire-storms, each consisting of physical skirmishes among individual combatants, but as a complicated interleaving network of physical action resulting from local interpretations of local information. The problem is to articulate what we mean by information.

Does the presence of fractals in combat point to something fundamental?

Are the fractal-like power-law scalings that have been observed for casualty rates and message traffic flow (see pages 74-75) indicative of some deep underlying process (akin to self-organized criticality) that can be exploited, or are they merely interesting curiosities and nothing else?

Appendix A: A Summary of the Mathematical Tools of Nonlinear Dynamics and Complex Systems Theory Discussed in Part I of this Report

Below is a partial summary of the mathematical tools of nonlinear dynamics and complex systems theory discussed in Part I [28] of this report:

- **Qualitative Characterization of Chaos.** Four qualitative methods for verifying the presence of chaos in a system were discussed. These included looking at the system's *time-dependent behavior*, using a *Poincare plot* to reduce the dimensionality, calculating the *autocorrelation function* and observing the *power spectrum* for the system.
- **Quantitative Characterization of Chaos.** Three sets of quantitative measures of chaos were introduced, including *Lyapunov exponents* (that measure the exponential divergence of initially nearby trajectories), *generalized fractal dimensions* (that, roughly speaking, measure the minimum number of variables needed to specify a chaotic attractor), and *the Kolmogorov-Sinai entropy* (that measures the rate of information gain per unit time in observing a chaotic system).
- **Time-Delayed Embedding.** The embedding technique is a method for reconstructing a state space from time-series data. It assumes that if the embedding dimension is large enough, the behavior of whatever system is responsible for generating the data can be described by a finite dimensional attractor. Its main strength lies in providing detailed information about the behavior of degrees-of-freedom other than the ones that are directly observed.
- **Chaotic Control.** Chaotic control exploits the fact that chaotic systems exhibit sensitivity to initial conditions to stabilize regular dynamical behaviors and thereby effectively "direct" chaotic trajectories to a desired state.

- **Cellular Automata.** Cellular automata are a class of spatially and temporally discrete, deterministic dynamical systems characterized by local interaction and an inherently parallel evolution. They serve as prototypical mathematical models of complex systems, and appear to capture many essential features of complex self-organizing cooperative behavior observed in real systems.
- **Genetic Algorithms.** Genetic algorithms are a class of heuristic search methods and computational models of adaptation and evolution based on natural selection. Genetic algorithms mimic and exploit the genetic dynamics underlying natural evolution to search for optimal solutions of general combinatorial optimization problems. This very powerful tool is used frequently as the backbone of many artificial life studies.
- **Agent-Based Simulations.** Agent-based simulations of complex adaptive systems are predicated on the idea that the global behavior of a complex system derives entirely from the low-level interactions among its constituent agents. By relating an individual constituent of a complex adaptive system to an agent, one can simulate a real system by an artificial world populated by interacting processes. Agent-based simulations are particularly adept at representing real-world systems composed of individuals that have a large space of complex decisions and/or behaviors to choose from.
- **Swarm.** *Swarm* (currently under development at the Santa Fe Institute) is a multi-agent simulation platform for the study of complex adaptive systems. The goal of the *Swarm* project is to provide the complex systems theory research community with a fully general-purpose artificial-life simulator. *Swarm* has been intentionally designed to include as few ad-hoc assumptions about the design of a complex system as possible, so as to provide a convenient, reliable and standardized set of software tools that can be tailored by researchers to specific systems. Another multi-agent model, called ECHO and developed by John Holland, is not as suitable for general purpose modeling as *Swarm* because of its many biases and in-built assumptions about the functioning of natural ecologies.

- **Neural Networks.** Neural nets represent a radical new approach to computational problem solving. Their *bottom-up* methodology stands in stark contrast to traditional *top-down* approach to artificial intelligence (AI). The approach is inspired by such basic skills of the human brain as its ability to continue functioning with noisy and/or incomplete information, its robustness or fault tolerance, its adaptability to changing environments by learning, etc. Neural nets attempt to mimic and exploit the parallel processing capability of the human brain in order to deal with precisely the kinds of problems that the human brain itself is well adapted for; in particular, pattern recognition.

Appendix B: Nonlinear Dynamics and Chaos

Key Concepts

Some of the key concepts of nonlinear dynamics and chaos include the following:

- **Phase Space.** Phase space is a mathematical space spanned by the parameters that describe a dynamical system's behavior. If the system is described by an ordinary differential flow the entire phase history is given by a smooth curve in phase space. Each point on this curve represents a particular state of the system at a particular time. For closed systems, no such curve can cross itself. If a phase history a given system returns to its initial condition in phase space, then the system is periodic and it will cycle through this closed curve for all time. For example, a mechanical oscillator moving in one-dimension has a two-dimensional phase space spanned by the position and momentum variables.
- **Nonlinear Feedback.** A key dynamical mechanism responsible for deterministic chaos is nonlinear recursive feedback (or *mixing*). Chaos is the result of a cascading series of feedback chains whereby one variable affects another which in turn affects the first, and so on. Think of a breadmaker's dough, as it is *stretched* and *folded* during kneading. A single lump of dough undergoes a seemingly erratic behavior as it is stretched and folded to produce a widely distributed, and finely structured, sheath that becomes intertwined with other stretched and folded sheaths of dough. The same basic mechanism is responsible for the appearance of the finely ordered microstructure immersed in what appears to be global disorder in chaotic systems.
- **Sensitivity to Initial Conditions.** Deterministic chaos is characterized chiefly by the so-called "Butterfly Effect," which alludes to the fact that two initially nearby points of a chaotic trajectory diverge exponentially in time.
- **Unpredictable Determinism.** Sensitivity to initial conditions implies that, despite the dynamics of a system being

rigorously deterministic, the long-term behavior of such a system appears irregular and is unpredictable.

- **Bifurcations.** A Bifurcation is the splitting into two modes of behavior of a system that previously displayed only one mode. It represents a transformation from one type of behavior into a qualitatively different type of behavior. This splitting occurs as a control parameter is continuously varied. In the logistic equation, for example, a period-doubling bifurcation occurs whenever all the points of a period- 2^n cycle simultaneously become unstable and the system becomes attracted to a new period 2^{n+1} cycle.
- **Fractals.** A fractal structure is characteristic of chaotic phenomena. Loosely speaking, fractals are geometric objects characterized by some form of self-similarity; that is, parts of a fractals, when magnified to an appropriate scale, appear similar to the whole. Finer and finer magnification reveals smaller and smaller versions of essentially the same structure on all levels. A more technical definition of fractals is that they are objects whose fractal (or *Hausdorff*) dimension does not equal its topological dimension.
- **Universality.** Universal behavior, when used to describe the behavior of a dynamical system, refers to behavior that is independent of the details of the system's dynamics. It is a term borrowed from thermodynamics. According to thermodynamics and statistical mechanics the critical exponents describing the divergence of certain physical measurables – such as specific heat, magnetization or correlation length – are universal at a phase transition in that they are essentially independent of the physical substance undergoing the phase transition and depend only on a few fundamental parameters (such as the dimension of the space). Similarly, in nonlinear dynamics, it has been shown that all equations of a particular universality class exhibit the same universal qualitative and quantitative patterns of behavior. Feigenbaum's universal convergence rate of (the single) control parameter for all one-dimensional maps that have a single quadratic maximum on the unit interval (see Part I, page []) is one example.
- **Strange Attractors.** A strange attractor is an attractor – which means it represents a region of phase space that the

system inevitably approaches as it evolves – that also displays sensitivity to initial conditions. This has the important consequence that while the behavior for each initial point may be accurately followed for short times, prediction of long-time behavior of trajectories lying on strange attractors becomes effectively impossible. However, strange attractors also reveal a long-term trend in the overall dynamics. That is, they reveal an underlying global pattern to motion that, locally, appears chaotic. Strange attractors also frequently have a self-similar or fractal structure.

- **Multiple Attractors.** Nonlinear systems typically harbor multiple basins of attraction and/or multiple attractors. Moreover, the boundaries between basins can have very complicated fractal forms. The manipulation of transitions among two or more attractors can potentially provide important insights for strategic options.

Basic Lessons

The major lesson of nonlinear dynamics and deterministic chaos is that a dynamical system does not have to be "complex" or to be described by a large set of equations, in order for the system to exhibit chaos – all that is needed is an embedded source of nonlinearity.

Chaos teaches us that even though many models are susceptible to seemingly erratic behavior, if that behavior stems from deterministic chaos then there is still the hope of identifying relevant trend and patterns if the behavior of the model's output is studied for a long enough period of time. In other words, apparently random output may be indicative of a long-term order that will become discernible only after the system has been observed for a long enough period of time and/or has been explored over a large enough region of its phase space.

Basic lessons of nonlinear dynamics and deterministic chaos include:

- **Chaos is Pervasive.** Nonlinear dynamics teaches us to appreciate the fact that chaos is pervasive. Indeed, because nonlinearity is such a common occurrence in nature, almost all real physical systems harbor chaotic behavior for some parameter regimes. Stanislaw Ulam once suggested

that "to call the study of chaos 'nonlinear science' is like calling zoology the study of non-elephant animals." (Gleick, p68 []). While linearity is a convenient approximation of reality, it fails to adequately describe most real-world systems. This pervasiveness of nonlinearity suggests that whenever "apparently random" behavior is observed in a real system, it can probably be described by deterministic (non-random) chaos.

- **Small Perturbations can Induce Large Changes.** Small changes in the control parameters of nonlinear systems can lead to major qualitative transitions of behavior. Contrast this with the more traditional ("linear") view which assumes that small perturbations lead to only small changes in a system's behavior. Furthermore, knowledge of how bifurcations arise in chaotic systems (1) tells us what kinds of transitions to expect when we add feedback to a system, and (2) suggests ways in which to *selectively adjust* feedback so as to induce desired transitions.
- **Behavior Depends on Location in Phase Space.** The kind of behavior a dynamical system exhibits – whether it is a simple limit cycle, a periodic orbit, or a strange attractor with fractal properties – depends on where the dynamical systems "lives" in its parameter space.
- **Universality Implies Predictability and Simplicity.** Universality (see above) has two very important implications: (1) *quantitative predictability*, whereby, for example, the spacing between and convergence rate between different regimes of dynamical behavior can be quantitatively predicted; and (2) a *simplification of complexity*, by asserting that large classes of systems all behave qualitatively and quantitatively the same way. "The wonderful thing about this universality is that it does not matter much how close our equations are to the ones chosen by nature; as long as the model is in the same universality class ... as the real system, both will undergo a period-doubling sequence. That means that we can get the right physics out of very crude models." Cvitanovic []
- **Chaos Contains Embedded Structure and Pattern.** While individual trajectories of a chaotic system appear to be erratic or even random, the attractors of the system harbor important information about certain recurrent aspects of its

long-term behavior. For example, while a strange attractor may not help us to predict the step-by-step evolution of a system along one of its possible trajectories, its structure tells us much about what the overall trajectory will look like. Moreover, the relative time that an orbit spends visiting various parts of a strange attractor gives us an idea about how likely the system is to be in certain parts of its phase space.

- **Chaos Does Not Preclude Short-Term Predictability.** Given sufficient data, techniques such as time-delayed embedding provide short-term predictions about a system's behavior, even if the system is chaotic. Moreover, these predictions can be made even when the underlying dynamics is not known. Lyapunov exponents quantify the limits of predictability.
- **Chaos can be "Controlled."** The extreme sensitivity of chaotic systems to small perturbations to initial conditions can be exploited to stabilize regular dynamic behaviors and to effectively "direct" chaotic trajectories to a desired state. Moreover, this can be done using only experimental data in which no model is available for the system. It is interesting to point out that this is a capability that has no counterpart in nonchaotic systems for the ironic reason that the trajectories in nonchaotic systems are stable and thus relatively impervious to desired control. A recent survey article [] lists applications for communications (in which chaotic fluctuations can be put to use to send controlled, pre-planned signals), for physiology (controlling chaos in heart rhythms), for fluid mechanics and chemical reactions. The important point to remember is that once a system has been shown to be chaotic, its attractor *must* contain a dense set of unstable periodic orbits and is therefore susceptible to some form of chaotic control.
- **Low-Dimensional versus High-Dimensional Chaos.** Chaotic dynamics is often misinterpreted to mean *random* dynamics. Strictly speaking, since chaos is spawned from a deterministic process, its apparent irregularity stems from an intrinsic magnification of an external uncertainty, such as that due to a measurement of initial conditions. Sensitivity to initial conditions amplifies an initially small uncertainty into an exponentially large one; or, in other words, short-term determinism evolves into long-term

randomness. Thus, the important distinction is not between chaos and randomness, but between chaotic dynamical systems that have low-dimensional attractors and those that have high-dimensional attractors. For example, if a time series of evolving states of a system is generated by a very high dimensional attractor (or if the dynamics is modeled in a state space whose dimension is less than that of the attractor), then it will be essentially impossible to gather enough information from the time series to exploit the underlying determinism. In this case, the apparent randomness will in fact have become a very real randomness, at least from a predictability standpoint. On the other hand, if the time series is generated by a relatively low dimensional attractor, it is possible to exploit the underlying determinism to predict certain aspects of the overall behavior. The *Information dimension* can be used to estimate the minimum number of variables needed to describe a system. Moreover, if a system can be shown to have a small non-integer dimension, it is probable that the underlying dynamics are due to nonlinearities and are not random.

Appendix C: Complex Systems Theory

Key Concepts

Some of the key concepts of complex systems theory include the following:

- **Autonomous Agents.** Autonomous agents are entities that, by sensing and acting upon their environment, try to fulfill a set of goals in a complex, dynamic environment. They can sense the environment through sensors and act on the environment through actuators; they have an internal information processing and decision making capability; and they can anticipate future states and possibilities, based on internal models (which are often incomplete and/or incorrect). Since a major component of an agent's environment consists of other agents, agents spend a great deal of their time adapting to the adaptation patterns of other agents.
- **Self-Organization.** Self-organization is a fundamental characteristic of complex systems. It refers to the emergence of macroscopic nonequilibrium organized structures due to the collective interactions of the constituents of a complex system as they react and adapt to their environment.
- **Non-equilibrium.** A system is said to be in equilibrium when it is in a particularly simple, quiescent state such that its properties are constant and spatially and temporally uniform. The most *uninteresting* systems, from the point of view of complex systems theory are systems that are in equilibrium. The most interesting systems are those that exist in far-from-equilibrium states, continually seeking new ways to adapt to their environment.
- **Co-adaptation and co-evolution.** Coadaptation refers to the mutually selective forces acting on entire groups of organisms in an ecology – or autonomous adaptive agents in an artificial life ecology – to accumulate favorably interacting genes in the gene pool of the population. Complex systems deal with not just one organism adapting to a given set of circumstances, but with many organisms, all

adapting to, and evolving with, all of the organisms that make up their environment.

- **Decentralized Order.** Decentralized order refers to the fact that the spontaneous appearance of order in a complex system is typically due solely to *parts* acting locally on local information. The global order thus emerges without any need for external control. There is no God-like "oracle" dictating what each and every part ought to be doing.
- **Phenotype and Genotype.** Phenotype refers to the observable characteristics and properties of an organism. Genotype refers to the actual genetic constitution of an organism. However, these basic definitions have a deeper metaphorical significance. They point to the dichotomy that exists between, on the one hand, the dynamics that defines micro-level of a complex system, and, on the other hand, the macroscopic behavioral properties of that system. Knowing an organisms genotype (or, say, the underlying lattice-rules for a cellular automaton), does not necessarily tell you anything at all about an organisms phenotype (or how the cellular automaton will actually behave). Understanding the nature of the connection the genotype level and phenotype-level remains one the deepest, most profound issues in complex systems theory today.
- **Self-Organized Criticality.** Self-organized criticality (SOC) describes a large body of both phenomenological and theoretical work having to do with a particular class of time-scale invariant and spatial-scale invariant phenomena. Fundamentally, SOC embodies the idea that dynamical systems with many degrees of freedom naturally self-organize into a critical state in which the same events that brought that critical state into being can occur in all sizes, with the sizes being distributed according to a power-law. The kinds of structures SOC seeks to describe the underlying mechanisms for look like equilibrium systems near critical phase-transition points but are not near equilibrium; instead, they continue interacting with their environment, "tuning themselves" to a point at which critical-like behavior appears. Introduced in 1988, SOC is arguably the *only* existing holistic mathematical theory of self-organization in complex systems, describing the behavior of many real systems in physics, biology and economics. It is also a universal theory in that it predicts

that the global properties of complex systems are independent of the microscopic details of their structure, and is therefore consistent with the "the whole is greater than the sum of its parts" approach to complex systems. Put in the simplest possible terms, SOC asserts that complexity is criticality. That is to say, that SOC is nature's way of driving everything towards a state of maximum complexity.

Basic Lessons

The major lesson of complex systems theory is that complex behavior is usually an emergent self-organized phenomenon built upon the aggregate behavior of very many nonlinearly interacting "simple" components. It advocates, in essence, a *simplicity breeds complexity* approach to the study of complex systems. There are myriad examples of ostensibly high-dimensional complex systems (that is, systems composed of a very large number of degrees-of-freedom) such that when "tuned" with an appropriate control parameter effectively behave as though they are low-dimensional. Think of the convective rolls of a fluid that is heated on the bottom and cooled at the top. Though the liquid itself consists of millions of interacting molecules, the motion of convective rolls – which can be simple and even chaotic – is well described by a *single* "roll amplitude" parameter.

Basic lessons of complex systems theory include:

- **Nonlinearity is Key.** Without nonlinear interactions there can be no deterministic chaos in simple systems and no complex behavior in complex systems. Moreover, nonlinear systems appear to be much more pervasive than linear systems. By virtue of nonlinearity, the behavior of the "whole" is not just a simple aggregate of the constituent "parts."
- **Interconnectivity is Important.** How the parts of a complex system are interconnected is just as important as what those parts are and what does parts do.
- **Parts have Meaning Only Within Context.** The effect that parts have on the remainder of the system – literally, how those parts are defined within the complex system – is determined by the context of the whole within which those

parts exist. In referring to any part P of a complex system, one must also point to various other parts with which P interacts (or may interact in the future).

- **"Process and Evolution" vice "Solution".** Simple (i.e. low-dimensional) dynamical systems are characterized by simple attractors – fixed points, limit cycles, quasiperiodic and chaotic (or strange) attractors. Although one can also try to characterize the behavior of complex systems with these attractor "labels," such a description would entirely miss the essence of what it means to be a complex system. A complex system embodies process, a ceaseless search for a better "solution" for an ill-defined, amorphous ever receding "problem." There is no such thing as "the solution," as the problem continually changes. In Zen-like fashion, you can say that the harder one tries to pin-down the behavior of a complex system with some static measure, the further one is from understanding what the complex system is really doing.
- **Adaptability.** The essence of a complex adaptive system is that its constituent parts are not Newtonian "billiards" that react blindly (but in well-defined fashion) to the world around them, but are instead endowed with an ability to *sense, learn from, and adapt* to their environment as they and the environment both evolve in time. An related lesson is that individual solutions (or evolutionary timelines) are essentially non-reproducible; a given system may "solve" a given problem in many different ways.
- **Emergence.** Perhaps the central concept of complex systems theory is that high-level behaviors emerge naturally out a brewing soup of low-level interactions. A flock of birds (or "Boids," see part I [28], page 73) does not need a central direction to behave in an apparently orchestrated manner. Nowhere on the lattice rule-level in Conway's *Life* CA game (see part I [28], page 87) is there any hint of the particle-like glider that spontaneously emerges on a higher level, and then apparently obey a dynamics all its own. The lesson is that where there is an assemblage of very many nonlinearly interacting parts, there is a good possibility of emergent behaviors on higher levels than those defining the underlying interactions. Moreover, such emergent behavior can appear on multiple spatial and temporal levels.

- **Global Order Arises from Local Activity.** Complex high-level patterns are often due to relatively simple local dynamics lying on a much lower level. The intricate swirls of the cellular automaton model of the Belousov-Zhabotinski reaction (see Part I [28], page 90), for example, have a characteristic length scale that is on the order of twenty or more individual lattice sites who local (i.e. *one site wide*) dynamics is responsible for those patterns.
- **Global order Affects Local Dynamics.** Not only does an emergent high-level structure generally owe its existence to low-level local dynamics, but the high-level patterns also affects the local behavior. Think of a vortex in a turbulent flow of liquid. On the one hand, the local interactions among the individual molecules making up the fluid are directly responsible for producing the vortex. On the other hand, once the vortex is formed it dictates the flow of molecules that surround it by letting in some and releasing others into the surrounding liquid.

Appendix D: Irreducible Semi-autonomous Adaptive Combat Agents

One obvious application of complex system theory to land warfare on Tier-VII of the eight tiers of applications is to model land combat as a co-evolving "ecology" of local-rule-based semi-autonomous agents. In this appendix we present a brief outline of the basic element of such a model, which we call an Irreducible Semi-Autonomous Adaptive Combat Agent (ISAACA).⁵⁰

An ISAACA represents a primitive combat unit (infantryman, tank, transport vehicle, etc.), that is equipped with:

- a default local rule set specifying how to act in a generic environment; i.e. embedded "doctrine"
 - acts may consist of
 - simple situational assessment
 - communication of information (up/down echelon)
 - movement
 - advance
 - retreat
 - obstacle avoidance
 - movement is both "individually motivated" – i.e. each ISAACA's primary instinct is survival – and locally driven – i.e. each ISAACA acts according to what its nearby friends are doing)
 - One possible "template" for goal-directed *movement* is to continually optimize the trade-offs among (1) choosing the fastest route from starting point to goal, (2) maintaining a minimal lethality of surrounding enemy

⁵⁰ Apart from its descriptive value, the acronym ISAACA was chosen to pay tongue-in-cheek homage to Isaac Newton. It seemed an appropriate choice to make given that the "new sciences" represents a fundamental shift *away* from linear "Newtonian" thinking.

ISAACAs, and (3) maintaining the maximal distance to all approaching enemy ISAACAs.

- combat
- a goal (or goals) directing its behavior
- an internal adaptive mechanism allowing it to alter its default rule set that acts according to an internal map of its environment
 - motivated by an internal value system (perhaps patterned after Smith's "Calculus of Ethics" [])
 - adaptation via GA: each ISAACA effectively "plays out a scenario" using genetically-encoded set of possible tactics; fitness is "expected payoff" modulo an internal value system
- a hierarchical rule set, consisting of orders passed down echelon via the C^2 topology
- a global rule set that determines combat attrition and reinforcement

What is the basic idea?

- to abstract the universal patterns of behavior and/or strategies/tactics that are essentially independent of the details of the makeup of individual ISAACAs
- to model land warfare by focusing on an ecology of ISAACAs, the decision process, and its role in the C^2 hierarchy:
 - the premise is that an ISAACA's internal processing and its ability to react and adapt to continually changing external stimuli is as important a dynamical driver as is the set of "firepower statistics" typically used in calculating force-on-force attrition
 - ISAACAs are semi-autonomous agents making local decisions predicated on rule-based doctrine codification – agent options are constructs of an agent's local perception (or mental map) of goal-vs-lethality

trade-off and are (simultaneously) constrained by global commands issued by higher echelon

- to embody objective and doctrine within each ISAACA, and endow each ISAACA with a sufficient “decisional flexibility/robustness” to continually adapt its default rule sets as (perceived) situation warrants
- to observe what global combat patterns *emerge* from the local agent- and rule- based skeletal substrate
- to augment a conventional CA approach in three ways:
 - *embed an intrinsic adaptability*
 - not just an evolution according to a fixed set of rules, but an evolution of the rules themselves (via rule templates)
 - *allow for non-local information* (via an embedded C^2 topology) and command hierarchy
 - *allow for GA-derived pseudo-global strategies*
 - what sequence of local strategies yields “optimal” results modulo prescribed goal? – orders pumped down echelon are based on “evolved” tactics played out on possibly imprecise mental maps of local/global commanders

Actions, Goals, Properties, and States

Given an ISAACA, say A:

- what are the variables defining A’s current state?
 - self-identity (“Red”, “Blue”)
 - rank (infantry, captain, colonel, general)
 - weapon store
 - maximum firepower
 - minimum firepower

- maximum range
- offensive capability: how good is A at hitting and killing an enemy unit that it fires at
 - aim accuracy
 - probability of hit
 - probability of kill (depends on enemy unit aimed at)
- defensive capability : how good is A at avoiding being hit and killed by an enemy unit
 - maneuverability
 - terrain camouflage
 - armor (enemy firepower required for given fractional change in overall “health”)
 - vulnerability to each kind of weapon
- movement
 - maximum range/time-step
 - direction
 - speed
 - terrain-dependence
- adaptive rule-set
 - internal rules determine all of A’s actions at time t , but also evolve in time according to changes in A’s environment
- morale
 - increases with (perception of) local combat “successes”)
 - decreases with increasing damage
 - morale may be used to define ISAACA “profiles,” describing different “personalities”
- combat “quality”

- higher quality assures higher attack “success” rate (i.e. firepower and maneuverability distributions weighted more towards higher end, say with a tail-end-weighted beta-distribution); in short, higher quality units “perform better”
- increases with A’s experience
- higher quality assures lesser degradation of morale under adverse conditions
- health
 - defines A’s overall health and functionality; when health = 0, A is “dead”
- what does A know about its environment? (what is A’s mental-map?)
 - terrain (i.e. lattice-cells) within visual line-of-site (LOS)
 - identity of ISAACAs within visual LOS
- what can A do?
 - *sense environment*: A has an associated LOS-radius R_{LOS} that defines an area within which it can sense terrain and other ISAACAs
 - *communicate*: with ISAACAs of next-higher rank
 - *move*: A can move anywhere within an associated movement-radius R_{move} , provided that it is not impeded by any intervening terrain and/or other constraints; where A moves is determined by A’s adaptive rule-set
 - engage and be engaged in combat
 - formulate local strategy
- with whom does A communicate?
 - communicates directly only with those ISAACAs that are next-higher in rank
 - A transmits its location and mental map

- A receives orders (overriding default and/or locally formulated tactics)
- what are A 's objectives?
 - *survival* (perhaps defined by a “maximal enemy firepower” threshold that enables a local-rule override – “retreat!” – over any default or higher-echelon orders)
 - achieve globally prescribed “victory” conditions
 - capture a prescribed area
 - destroy enemy “headquarters”
 - protect friendly “headquarters”
 - maintain prescribed set of friendly “structures”
 - maintain prescribed fraction of friendly forces
 - kill a prescribed number of enemy ISAACAs or kill prescribed fraction of enemy forces
- what are A's actions and strategies predicated on? (how are A's objectives defined?)
- how does A adapt to its environment?
 - each A chooses a strategy by effectively playing out a series of internalized “games” predicated on an internal map of the combat “playing field”
 - A uses the strategy with the highest perceived “payoff” (which must also be consistent with any pertinent higher-echelon orders)
 - “strategy” involves both movement and combat
- how does A distinguish between friendly and enemy forces?
 - relative positioning?
- what are the different kinds of engagements?
 - *direct* A \leftrightarrow A' *combat*: A “sees” A', and vice-versa, and both engage in one-on-one combat; outcome is determined probabilistically, taking into account

“weapon strength,” range, range, morale, defender’s strength and visibility

- $A \leftrightarrow Area(A')$: A “knows” or “suspects” A' is located within an area $Area(A')$ – consisting of, say, an N-by-N array of lattice cells – and “blindly” fires at a random cell or cells in $Area(A')$
- *collective Fire*, $F(R,A)$: A coordinates its fire into a patch of enemy territory – an area of size A and at range R –with nearby As
- how are many-vs-many – i.e. N_R “Red” As simultaneously come within LOS of N_B “Blue” As – conflicts resolved?
 - *independent fire* – when all ISAACAs act independently (i.e. when units from one side can see all units of the other side but none from their own):
 - *mixed fire* – when one side acts independently and the other collectively (i.e. when each unit on one side can see all opposing units but none from its own, and units from the opposing side can see both sides):
 - *collective fire* – when all ISAACAs act collectively (i.e. when units on both sides are aware of all other units):
- what is A’s role in the C^2 hierarchy?
- what combat doctrine can be captured by local rules?
- what orders can passed down echelon to A?
 - reconnaissance
- what are the different kinds of terrain?
 - flat/rough
 - road
 - forest
 - hill

- river
- minefields
- structures (bridges, storage, bunkers, etc.)
- how is terrain characterized?
 - altitude (affecting LOS)
 - movement (“passability”) index
 - function of type of ISAACA occupying cell
 - camouflage (“fog”) index (affects visibility and/or identifiability)
 - *visibility*: ISAACAs positioned in a forest cell *C* are visible only to ISAACAs immediately adjacent to *C*
 - *defensive capability*: fractional increase of defensive-capability allotted to all ISAACAs within cell of given terrain-type
- how does A interact with different kinds of terrain?

Decision Dynamics

Each ISAACA:

- thinks for itself
 - acts according to default “rules” in order to bring it closer to achieving its general goal
- follows local group action
- engages/is-engaged-by enemy fire
- obeys (non-local) orders issued from higher echelon
 - follows orders from local commander (on time scale)
 - obeys global commander via its local commander (on time scale)

- each ISAACA “sees” only what is within its LOS-radius
 - each ISAACA generates (via an embedded GA algorithm) a set of local movement and fire strategies
- each local commander (LC-ISAACA) is aware of each of the local environments as seen by the ISAACAs under its command (in addition to seeing what is within its own LOS-radius)
 - each LC-ISAACA generates a set of movement and fire-conflict strategies for each of the ISAACAs in its command
- each global commander (GC-ISAACA) has a (possibly inaccurate) picture of the overall combat environment, and is an aggregate of environmental perceptions of the local commanders under its command
 - each GC-ISAACA generates a set of movement and conflict strategies for each of the LC-ISAACAs in its command

Measures of Combat effectiveness

- local
 - sustainability
 - rate of advance
- global
 - total relative attrition (enemy to friendly casualties)
 - attrition rate inflicted on enemy
 - time to reach goal

Appendix E: A Brief Discussion of the General Applicability of Genetic Algorithms

Even a brief glance at the list of possible applications of complex systems theory to land warfare (see, for example, table 3), shows that genetic algorithms figure prominently on many of the eight tiers of applicability. The fact that it is so testifies to the enormous popularity as a basic research tool of this computational technique in complex systems theory. Indeed, there are few artificial-life computer models that do not include at least some form of the basic genetic algorithm.

But, given that there are many flavors of genetic algorithms, how does one know, a-priori, if a particular realization of a genetic algorithm is the best one to use? How do we that other methods to solve combinatorial optimization problems – such as gradient descent and simulated annealing – do not perform better than genetic algorithms for a given instance of a problem? These questions are actually difficult, if not impossible, to answer, and a one has to be very careful in applying genetic algorithms. In this appendix we briefly discuss the general applicability of genetic algorithms, and suggest that the algorithm should not be considered a panacea solution to any problem.

Genetic algorithms are, in fact, extremely powerful heuristic tools for finding near-optimal solutions for general combinatorial optimization search problems. They have been successfully applied to a wide-variety of problems, and not just those confined to complex systems theoretic studies. Problems they have been used for include traveling salesman problems, VLSI circuit layout, gas pipeline control, the parametric design of aircraft, neural net architecture, models of international security, strategy formulation, and so on.

Recall (see Part I [28], pages 93-101) that genetic algorithms consist essentially of five basic steps:

- *Step 1:* begin with a randomly generated population of chromosome-encoded **templates** of solutions to a given problem

- *Step 2:* calculate the fitness of each chromosome, where fitness is a measure of how well a member of the population performs at solving the problem
- *Step 3:* retain only the fittest members and discard the least fit members
- *Step 4:* generate a new population of chromosomes from the remaining members of the old population by applying operations such as reproduction, crossover, and mutation
- *Step 5:* calculate the fitness of these new members of the population, retain the fittest, discard the least fit, and re-iterate the process

The two key words in the above list are "templates" and "fitness." Genetic algorithms are generally applicable whenever one is faced with a problem whose "solution" is known to reside (or suspected of residing) somewhere within a possibly vast N-dimensional "solution-space," but its exact location in that space is unknown. However, because a genetic algorithm generally needs a fitness function – that is, a measure of how close a candidate solution is to the desired solution - to conduct its search of the solution space, great care must be put into defining the fitness measure.

If the fitness is either poorly defined, or, worse, leaves out the relevant parameters of the problem altogether, the genetic algorithm will effectively begin its search in a "brain-dead" state, and will likely fail in its attempt. While open-ended genetic algorithms, that start out having *no fitness function*, are often used in artificial-life studies of the evolution of natural ecologies, the reader is reminded that natural ecologies differ fundamentally from a natural "combat ecology." While natural ecologies tend to evolve their own (changing) fitness functions over time, a combat environment usually comes pre-defined, with a litany of constraints and measures. While that is not to say that combat ecologies do not also generate their own fitness measures in the course of their evolution, it is still true that any reasonable model of such ecologies would be remiss if it did not respect the pertinent control parameters defining them. Therefore, for typical combat-related problems, one must think carefully about what fitness measure is relevant for what parameters *before* a genetic algorithm is used to explore the resulting solution space.

A genetic algorithm is only as good as the fitness measure it uses to explore its "solution space" with.

Keep in mind also that a genetic algorithm has a chance to come close to the desired "solution" *only* if (1) the template of the general solution has been defined accurately, so that it is known to reside somewhere in the overall solution space, (2) the fitness of a solution template has been defined correctly, so that the genetic algorithm knows, at all times, how close it is getting to the desired solution, and (3) the problem is, in principle, amenable to a genetic algorithm search. It is often impossible to satisfy all three conditions. Moreover, it is currently unknown how to decide, a-priori, how hard a particular fitness landscape will be for a genetic algorithm to search it for a solution. Some landscapes are more amenable than others for a genetic algorithm search. Of course, the even more general problem of determining what landscape is best suited for what general search technique remains unsolved.

The point here is that extreme caution is urged when applying genetic algorithms. One should not casually label a land warfare problem, that has been "solved" by using a genetic algorithm, a successful example of applying complex systems theory to land warfare solely on the basis of having used a genetic algorithm. Applying complex systems theory means thinking carefully about all of the different facets of what makes given system a "complex systems," and applying that knowledge intelligently and consistently to trying to understand what that system is really doing. A properly coded genetic algorithm that uses a poorly thought-out fitness measure for an even more poorly thought-out description of combat dynamics, is not just bad complex systems theory, it is bad science.

References

1. H. D. I. Abarbanel, *Analysis of Observed Chaotic Data*, Springer-Verlag, 1996.
2. *American Heritage CD-ROM Electronic Dictionary*, 3rd edition, Softkey International, 1994.
3. S. Banks, "Exploratory modeling for policy analysis," *Operations Research*, Volume 41, No. 3, May-June 1993, 435-449.
4. M. F. Barnsley and A. Sloan, "A better way to compress images," *Byte*, January 1988, 215-223.
5. M. F. Barnsley, *Fractals Everywhere*, Academic Press, 1988.
6. R. Beaumont, *War, Chaos, and History*, Wesport, Connecticut: Praeger, 1994.
7. R. K. Belew and M. Mitchell, editors, *Adaptive Individuals in Evolving Populations: Models and Algorithms*, Addison-Wesley, 1996.
8. A. Beyerchen, "Clausewitz, Nonlinearity, and the Unpredictability of War," *International Security*, volume 17, No. 3, Winter 1992/93, 59-90.
9. Bianco and Read, 1990, US patent number 5,048,086.
10. Col John Boyd, USAF (Ret.), *A Discourse on Winning and Losing*, August 1987.
11. R. Brooks, "Elephants don't play chess," pages 3-15 in *Designing Autonomous Agents*, edited by P. Maes, MIT Press, 1994.
12. M. E. Carter, C. Dean, C. Pollatz, A. Brink, D. Canody, P. Barry, T. Wallace, J. Dockery, and D. Littman, "Architecture for intelligent virtual agents that employ multiple learning strategies: or smart tanks avoid black holes," pages 263-270 in *Proceedings of the 3rd Conference on Computer Generated Forces and Behavioral Representation*, March 17-19, 1993, Orlando, Florida, Institute for

References

- Simulation and Training, Technical Report IST-TR-93-07, 1993.
13. F. Carvalho-Rodrigues, "A proposed entropy measure for assessing combat degradation," *Jour. Opl. Res. Soc. (UK)*, Volume 40, No. 8, 1989, 789-793.
14. J. Casti, *Reality Rules*, Volumes I and II, John Wiley and Sons, 1992.
15. J. Cooper, "Dominant battlespace awareness and future warfare," pages 103-120 in *Dominant Battlespace Knowledge: The Winning Edge*, edited by S. Johnson and M. Libicki, National Defense University Press, 1995.
16. E. Denning, *Cryptography and Data Security*, Addison-Wesley, 1982
17. P. J. Denning, "Modeling Reality," *American Scientist*, November-December 1990, 495-498.
18. J. T. Dockery and A. E. R. Woodcock, *The Military Landscape: Mathematical Models of Combat*, Cambridge, England: Woodhead Publishing Limited, 1993.
19. T. N. Dupuy, *Understanding War: History and Theory of Combat*, Paragon House Publishers, 1987.
20. C. Emmeche and J. Hoffmeyer, "From language to nature: the semiotic metaphor in biology," *Semiotica*, Volume 84, No. 1/2, 1991, 1-42.
21. J. M. Epstein, "On the mathematical biology of arms races, wars, and revolutions," pages 425-436 in *1991 Lectures in Complex Systems*, edited by L. Nadel and D. Stein, Addison-Wesley, 1992.
22. J. M. Epstein, "The adaptive dynamic model of combat," pages 437-459 in *1991 Lectures in Complex Systems*, edited by L. Nadel and D. Stein, Addison-Wesley, 1992.
23. S. Eubank and D. Farmer, "An introduction to chaos and randomness," pages 75-190 in *1989 Lectures in Complex Systems: The Proceedings of the 1989 Complex Systems Summer*

References

- School, Santa Fe, New Mexico, June 1989*, edited by E. Jen, Addison-Wesley, 1990.
24. H. A. Gutowitz, "Cryptography with dynamical systems," pages 237-274 in *Cellular Automata and Cooperative Systems*, edited by N. Boccara, E. Goles, S. Martinez, and P. Picco, NATO ASI Series Volume 396, Kluwer Academic Publishers, 1993.
 25. W. O. Hedgepeth, "The Science of Complexity for Military Operations Research," *Phalanx*, volume 26, No. 1, 1993, 25-26.
 26. W. D. Hillis, "Co-evolving parasites improve simulated evolution as an optimization procedure," *Physica D*, Volume 42, 1990, 228-234.
 27. A. Hubler, "Modeling and control of complex systems: paradigms and applications" pages 5-65 in *Modeling Complex Phenomena*, edited by L. Lam and V. Naroditsky, Springer-Verlag, 1992.
 28. A. Ilachinski, *Land Warfare and Complexity, Part I: Mathematical Background and Technical Sourcebook* (U), Center for Naval Analyses Information Manual CIM-461, March 1996, Unclassified.
 29. Jumarie, G., "A relativistic information theory model for general systems Lorentz transformation of organizability and structural entropy," *International Journal of Systems Science*, Volume 6, 1975, 865-886.
 30. S. Kauffman, *Origins of Order: Self-Organization and Selection in Evolution*, Oxford University Press, 1993.
 31. K. Kelly, *Out of Control: The Rise of Neo-Biological Civilization*, Addison-Wesley, 1994.
 32. J. A. S. Kelso, *Dynamic Patterns: The Self-Organization of Brain and Behavior*, MIT Press, 1995.
 33. J. T. Klein, *Interdisciplinarity: History, Theory and Practice*, Wayne State University: Detroit, 1990.

References

34. G. Lakoff and M. Johnson, *Metaphors we Live By*, University of Chicago Press, 1980.
35. R.C Lewontin, *The Genetic Basis of Evolutionary Change*, Columbia University Press, 1974.
36. V. Lesser, editor, *Proceedings of the First International Conference on Multi-Agent Systems*, MIT Press, 1995.
37. R. Lewin, *Complexity: Life at the Edge of Chaos*, Macmillan Publishing Company, 1992.
38. M. Libicki, "Mesh and the net: speculations on armed conflicts in an age of free silicon," McNair Paper 28, National Defense University, 1995. An electronic copy can be obtained from the WWW address: <http://www.ndu.edu/ndu/inss/macnair/mcnair28/m028cont.html>.
39. G. Mayer-Kress, "Nonlinear Dynamics and Chaos in Arms Race Models," pages 153-183 in *Modeling Complex Phenomena, Proceedings of the Third Woodward Conference, San Jose State University, April 12-13, 1991*, edited by Lui Lam and Vladimir Naroditsky, Springer-Verlag, 1992.
40. J. H. Miller, "Active nonlinear tests (ANTs) of complex simulation models," 5 Feb, 1996 (electronic postscript copy downloaded from <http://www.zia.hss.cmu.edu>.)
41. L. D. Miller, M. F. Sulcoski, *Foreign Physics Research with Military Significance: Military Implications of Research on Complex Adaptive Systems* (U), Defense S&T Intelligence Study, DST-1820S-440-93, December 1993 (Secret/NOFORN/ WNINTEL).
42. L. D. Miller, M. F. Sulcoski, *Foreign Physics Research with Military Significance: Applying Nonlinear Science to Military Programs* (U), Defense Intelligence Reference Document, NGIC-1823-614-95, March 1995 (Secret/ NOFORN/ WNINTEL).
43. M. Mitchell, *An Introduction to Genetic Algorithms*, MIT Press, 1996.
44. M. O'Hare, "Chaos pitch," *New Scientist*, 8 June 1996, 24-28.

References

45. A. Ortony, editor, *Metaphor and Thought*, Second Edition, Cambridge University Press, 1993.
46. E. M. Palmer, *Graphical Evolution: An Introduction to the Theory of Random Graphs*, John-Wiley and Sons, 1985.
47. L. M. Pecora and T. L. Carroll, "Synchronization in chaotic systems," *Physical Review Letters*, Volume 64, 1990, 821-824.
48. G. Piatetsky-Shapiro and W. J. Frawley, editors, *Knowledge Discovery in Databases*, MIT Press, 1991.
49. W. J. Ray, "EEG and chaos: Explorations in using dimensional estimation to estimate psychological state", in *Biobehavioral self-regulation in the East and West*. edited by Kikuchi, Springer-Verlag, Tokyo, 1993.
50. F. C. Richards, T. P. Meyer and N. H. Packard, "Extracting cellular automata rules directly from experimental data," *Physica D*, Volume 45, 1990, 189-202.
51. L. F. Richardson, *Statistics of Deadly Quarrels*, Pitsburg, PA: Boxwood, 1960.
52. J. Roughgarden, A. Bergman, S. Shafir and C. Taylor, "Adaptive computation in ecology and evolution: a guide for future research," Chapter 2 in *Adaptive Individuals in Evolving Populations: Models and Algorithms*, edited by R. K. Belew and M. Mitchell, Addison-Wesley, 1996.
53. A. Saperstein, "War and Chaos," *American Scientist*, Volume 83, November-December, 1995, 548-557.
54. J. F. Schmitt, *Chaos, Complexity and War: What the New Nonlinear Dynamical Sciences May Tell us About Armed Conflict*, Draft, 4 September 1995.
55. M. Schroeder, *Fractals, Chaos, Power Laws: Minutes from an Infinite Paradise*, W. H. Freeman and Company, 1991.
56. L. A. Segel, "Grappling with complexity," *Complexity*, Volume 1, No. 2, 1995, 18-25.

References

57. P. Senge, *The Fifth Discipline: The Art and Practice of the Learning Organization*, Currency Doubleday, 1990.
58. N. M. Smith, "A calculus for ethics: a theory of the structure of value. Part I," *Behavioral Science*, Volume 1, No. 2, April, 1956, 111-142.
59. N. M. Smith, "A calculus for ethics: a theory of the structure of value. Part II," *Behavioral Science*, Volume 1, No. 3, July, 1956, 186-211.
60. L. Steels and R. Brooks, editors, *The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents*, Lawrence Erlbaum Associates, Publishers: Hillsdale, New Jersey, 1995.
61. T. Tagarev, M. Dolgov, D. Nicholls, R. C. Franklin and P. Axup, *Chaos in War: Is it Present and What Does it Mean?*, Submitted to the Air Command and Staff College in fulfillment of the requirement of the AY94 Research Program, June 1994.
62. James G. Taylor, *Force-On-Force Attrition Modeling*, Operations Research Society of America, 1980.
63. James G. Taylor, *Lanchester Models of Warfare*, Volumes I and II, Operations Research Society of America, 1983.
64. Sun Tzu, *The Art of War*, translated by Thomas Cleary, Shambala, 1991.
65. L. E. Virr, A. Fairley and F. Yates, "Tactical application of genetic algorithms," *J. N. S.*, Volume 19, No. 3, 159-169.
66. M. Waldrop, *Complexity: The Emerging Science at the Edge of Order and Chaos*, New York: Simon and Schuster, 1992.
67. Lt. Col. I. G. R. Wing, "Chaos theory and intelligence analysis," *Australian Defence Force Journal*, Volume 115, Nov/Dec 1995, 22-26.
68. S. Wolfram, "Random sequence generation by cellular automata," *Advances in Applied Mathematics*, Volume 7, 1984, 123.